arXiv:1511.00255v1 [q-bio.NC] 1 Nov 2015

Neural ring homomorphisms and maps between neural codes

Carina Curto & Nora Youngs

November 3, 2015

Abstract

Understanding how the brain stores and processes information is central to mathematical neuroscience. Neural data is often represented as a neural code: a set of binary firing patterns $\mathcal{C} \subset \{0,1\}^n$. We have previously introduced the neural ring, an algebraic object which encodes combinatorial information, in order to analyze the structure of neural codes. We now relate maps between neural codes to notions of homomorphism between the corresponding neural rings. Using three natural operations on neural codes (permutation, inclusion, deletion) as motivation, we search for a restricted class of homomorphisms which correspond to these natural operations. We choose the framework of linear-monomial module homomorphisms, and find that the class of associated code maps neatly captures these three operations, and necessarily includes two others - repetition and adding trivial neurons - which are also meaningful in a neural coding context.

Contents

1 Introduction and Summary of Results **2**
- 1.1 Introduction . . . 2
- 1.2 Neural rings and the pullback map . . . 3
- 1.3 Neural rings as modules . . . 4
- 1.4 Compatible ring homomorphisms . . . 5
- 1.5 Neural ring homomorphism . . . 6
- 1.6 Examples . . . 7
- 1.7 Effect on the canonical form . . . 8

2 Homomorphisms and code maps **10**
- 2.1 Neural rings as modules . . . 12
- 2.2 Modules under different rings . . . 15

3 Linear-monomial module homomorphisms & neural ring homomorphisms **19**
- 3.1 Neural ring homomorphisms & proof of Theorem 1.4 . . . 21
- 3.2 Examples of neural ring homomorphisms . . . 23

4 Neural ring homomorphism and the canonical form **23**
- 4.1 Iterative algorithm for computing $CF(J_{\mathcal{C}})$ . . . 25
- 4.2 Proof of Algorithm 1 . . . 26

1 Introduction and Summary of Results

1.1 Introduction

A major challenge of mathematical neuroscience is to determine how the brain processes and stores information. By recording the spiking from a population of neurons, we obtain data about their relative activity. Much can be learned by considering this neural data in the form of a neural code: a set of binary firing patterns recorded from brain activity. In this context, a *neural code* on n neurons is a subset $\mathcal{C} \subset \{0,1\}^n$, with each binary vector in $\mathcal{C}$ representing a pattern of activity. This type of neural code is referred to in the literature as a combinatorial neural code [1] as it contains only the combinatorial information of which neurons fire together, without precise spike times or firing rates. These codes can be analyzed to determine important features of the neural data, using tools from coding theory [2] or topology [3]. In [4], we introduced the neural ideal and the corresponding neural ring, algebraic objects associated to a neural code that store its combinatorial information.

Thus far, work involving the neural ring has been primarily concerned with using the algebraic properties of the ring, such as the *canonical form* for the neural ideal, to extract structural information about the code [4, 5]. However, a neural code $\mathcal{C}$ is not an isolated object, but rather a member of a family of similar codes which may be obtained from $\mathcal{C}$ through certain natural operations, or maps. We define a *code map* from a code $\mathcal{C}$ to another code $\mathcal{D}$ to be any well-defined function $q : \mathcal{C} \to \mathcal{D}$. Considering the relationships between codes which are described by these maps allows us to determine precisely how information changes from one code to another. We are interested primarily in a set of 'nice' code maps which reflect a meaningful relationship between the corresponding neural data. Our primary motivating examples of 'nice' maps are listed below.

- **Permutation.** If $\sigma : [n] \to [n]$ is a permutation of $[n]$, then σ induces a permutation map q on any code $\mathcal{C} \subset \{0,1\}^n$, defined by $q((c_1, ..., c_n)) = (c_{\sigma(1)}, ..., c_{\sigma(n)})$ for any codeword $c = (c_1, ..., c_n) \in \mathcal{C}$. Applying a permutation to a code results in a code which is structurally identical to the original. That is, the labeling of the neurons is not important to the combinatorial properties of the code.

- **Deleting neurons.** Let $\alpha \subset [n]$ be a subset of neurons which we want to ignore. The deletion of these neurons is equivalent to projection onto the remaining neurons. That is, order $([n] \backslash \alpha) = \{i_1, ..., i_k\}$ so $i_1 < \cdots < i_k$; the code map deleting neurons in α is then given by $q((c_1, ..., c_n)) = (c_{i_1}, ..., c_{i_k})$ for any codeword $c = (c_1, ..., c_n) \in \mathcal{C}$.

- **Inclusion**: When two codes $\mathcal{C}$ and $\mathcal{D}$ on the same set of neurons are related by containment, then they are related by an inclusion map. That is, when $\mathcal{C} \subseteq \mathcal{D}$, the inclusion map is defined so $q(c) = c$ for any codeword $c \in \mathcal{C}$.

One type of code operation which we would **not** consider to be meaningful in the neural context is **bit-reversal**, where we map each binary vector to the vector with complementary support (that is, the map defined by $q((c_1, ..., c_n)) = (1 - c_1, ..., 1 - c_n)$). As an operation on neural data, it makes little sense to associate the activity of a subset of neurons to the activity of the complementary set.

Understanding the relationships between codes, like those described above, requires an understanding of how combinatorial structure changes when these maps are applied. Since the neural ring stores structural information, we here consider how operations on neural codes relate to algebraic operations between neural rings. Our question, therefore, is the following:

Question: What types of maps on neural rings correspond to 'nice' maps between codes?

The class of 'nice' maps certainly includes the three maps described above. In determining what other maps are in this class, we will in part allow the algebra to guide us. That is, the algebraic structure we choose for the best notion of *neural ring homomorphism* in order to capture these three maps may unavoidably capture other code maps. We are willing to consider those code maps 'nice' as well, provided they have a meaningful interpretation as an operation on neural data.

Describing relationships between maps on varieties and maps on their associated rings is not an uncommon consideration in algebraic geometry; see for example [6]. However, our goal is twofold: we wish to not only relate maps between codes to a notion of homomorphism between rings, but also to restrict algebraically to a special class of homomorphisms such that the associated code maps are similar to the maps described above. Finally, since permutation maps preserve all combinatorial structure of the code, we aim to define isomorphism between neural rings in some way which captures permutation maps.

The organization of this paper is as follows. In the remainder of Section 1, we outline our main results and explain our final choice for a notion of *neural ring homomorphism* which meets our stated restrictions. In Section 2, we provide proofs of the initial results about the correspondence between code maps and ring homomorphisms, and examine the structure of the neural rings as modules. In Section 3 we describe our choice of neural ring homomorphism which captures the desired features, and present the relevant proofs. Finally, in Section 4, we discuss the effect of our chosen code maps on the canonical form of the neural ideal, and exhibit an iterative algorithm for adapting the canonical form when a codeword is added.

1.2 Neural rings and the pullback map

First, we briefly review the definition of a neural code and its associated neural ring, and present the most immediate relationship between the set of functions from one code to another, and the set of homomorphisms between the corresponding neural rings.

Definition. A *neural code* on n neurons is a set of binary firing patterns $\mathcal{C} \subset \{0,1\}^n$. If $\mathcal{C} \subset \{0,1\}^n$ and $\mathcal{D} \subset \{0,1\}^m$ are neural codes, a *code map* is any assignment $q : \mathcal{C} \to \mathcal{D}$ sending each codeword $c \in \mathcal{C}$ to an image $q(c) \in \mathcal{D}$.

Given a neural code $\mathcal{C} \subset \{0,1\}^n$, we define the associated ideal $I_\mathcal{C} \subset \mathbb{F}_2[x_1, ..., x_n]$ as follows:

$$I_\mathcal{C} \stackrel{\text{def}}{=} \{f \in \mathbb{F}_2[x_1, ..., x_n] \mid f(c) = 0 \text{ for all } c \in \mathcal{C}\}.$$

The *neural ring* $R_\mathcal{C}$ is then defined to be $R_\mathcal{C} = \mathbb{F}_2[x_1, ..., x_n]/I_\mathcal{C}$, the ring of functions $\mathcal{C} \to \{0,1\}$.

Since the ideal $I_\mathcal{C}$ is precisely the set of polynomials which vanish on $\mathcal{C}$, we can make use of the ideal-variety correspondence and obtain a relationship between the code maps and ring homomorphisms immediately by using the pullback. Given a code map $q : \mathcal{C} \to \mathcal{D}$, each $f \in R_\mathcal{D}$ is a function $f : \mathcal{D} \to \{0,1\}$, and therefore we may "pull back" f by q to a function $f \circ q : \mathcal{C} \to \{0,1\}$, which is an element of $R_\mathcal{C}$. Hence for any $q : \mathcal{C} \to \mathcal{D}$ we may define the pullback map $q^* : R_\mathcal{D} \to R_\mathcal{C}$, where $q^*(f) = f \circ q$. That is, for every $f \in R_\mathcal{D}$, the following diagram commutes:

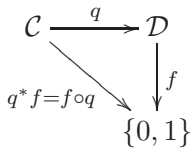

In fact, the pullback provides a bijection between code maps and ring homomorphisms, as the following theorem states.

Theorem 1.1. *There is a 1-1 correspondence between code maps $q : \mathcal{C} \to \mathcal{D}$ and ring homomorphisms $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$, given by the pullback map. That is, given a code map $q : \mathcal{C} \to \mathcal{D}$, its pullback $q^* : R_{\mathcal{D}} \to R_{\mathcal{C}}$ is a ring homomorphism; conversely, given a ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ there is a unique code map $q_\phi : \mathcal{C} \to \mathcal{D}$ such that $q^*_\phi = \phi$.*

Theorem 1.1 is proven in Section 2, and is a special case of Proposition 8, p.243 in [6]. The bijective correspondence described by the pullback is unsatisfying in its generality: any code map we can define has a corresponding ring homomorphism. Additionally, using the pullback we can show that two neural rings are isomorphic if and only if their corresponding codes have the same number of codewords (Lemma 2.3); no other similarity of structure is reflected by isomorphism. Therefore, we will go beyond the ring structure to determine an algebraic framework by which the activity of the neurons (and thus the combinatorial structure of the code) can be tracked.

1.3 Neural rings as modules

The neural ring is constructed by associating the activity of neuron i to the behavior of the variable x_i. When these rings are viewed abstractly without the details of their presentation (as with the pullback map), this relationship is lost. To ensure that we retain the presentation of the ring, we view the neural rings as modules under a ring on the variables x_i with no code-related relationships between the variables. This allows us to track each variable x_i via the module action, and hence to track the activity of neuron i.

For a neural ring $R_{\mathcal{C}}$ where $\mathcal{C}$ is a code on n neurons, we consider $R_{\mathcal{C}}$ as an $R[n]$-module, where $R[n] = \mathbb{F}_2[x_1, ..., x_n]/\langle x_i^2 - x_i \mid i = 1, ..., n\rangle$ is the neural ring on the full code $\{0,1\}^n$. The module action is standard multiplication of functions; see Section 2.1 for a precise definition. In particular, we note that $R_{\mathcal{C}}$ contains a basis of characteristic functions ρ_c, where

$$\rho_c(v) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } v = c \\ 0 & \text{if } v \neq c \end{cases}$$

The module action is then more easily described by noting that each variable $x_i \in R[n]$ either preserves a basis element or neutralizes it, depending on the associated codeword:

$$x_i \cdot \rho_c = \begin{cases} \rho_c & \text{if } c_i = 1 \\ 0 & \text{if } c_i = 0 \end{cases} = c_i \rho_c$$

For more about this module action, see Section 2.1. Most importantly, this module structure is not only a property of neural rings, but can be used to characterize them completely.

Theorem 1.2. *Suppose M is an $R[n]$-module. Then M is isomorphic to a neural ring $R_{\mathcal{C}}$ as an $R[n]$-module if and only if M has an $\mathbb{F}_2$-basis $\rho_1, ...\rho_d$ such that the following hold:*

1. *For all $i \in [n]$ and $j \in [d]$, $x_i \cdot \rho_j \in \{\rho_j, 0\}$.*
2. *For all $j, k \in [d]$ there exists at least one $i \in [n]$ such that exactly one of $x_i \cdot \rho_j$ and $x_i \cdot \rho_k$ is 0*

For such an $R[n]$-module, we can determine $\mathcal{C}$ from the module action.

Theorem 1.2 is proven in Section 2.1. This theorem shows that by considering the action of the variables x_i on a neural ring R, we can determine the unique code $\mathcal{C} \subset \{0,1\}^n$ for which $R = R_{\mathcal{C}}$. Having described a module framework which retains the information about the presentation of the neural ring, we consider how module homomorphisms relate to code maps. However, we need to do a bit of work first, since if we take our module structure as described above with no modification, we

would be unable to consider any relationship between neural rings on different numbers of neurons, as they are modules under different rings. In fact, the only code maps this structure would cover are inclusion maps between codes of the same length; this is the substance of Lemma 2.4. To address this issue, we turn to some well-known algebraic machinery that allows us to extend the notion of module homomorphisms to modules under different rings.

1.4 Compatible ring homomorphisms

Suppose R, S are rings with $\tau : R \to S$ a ring homomorphism. Any S-module M may then be viewed as an R-module via the homomorphism τ, using the action $r \cdot m = \tau(r) \cdot m$ for any $r \in R, m \in M$. In the neural ring framework, this says that given a ring homomorphism $\tau : R[m] \to R[n]$, we can consider the $R[n]$-module $R_{\mathcal{C}}$ as an $R[m]$-module. This motivates the following definition:

Definition. Given an R-module M, an S-module N, and a group homomorphism $\phi : M \to N$ (so $\phi(x + y) = \phi(x) + \phi(y)$), we say that a ring homomorphism $\tau : R \to S$ is *compatible* with ϕ if ϕ is an R-module homomorphism, where N is viewed as an R-module via τ. That is, τ is compatible with ϕ if $\phi(r \cdot x) = \tau(r) \cdot \phi(x)$ for all $r \in R, x \in M$. Equivalently, τ is compatible with ϕ if for every $r \in R$, the following diagram commutes:

$$\begin{array}{ccc} M & \xrightarrow{r\cdot} & M \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi} \\ N & \xrightarrow{\tau(r)\cdot} & N \end{array}$$

Having defined compatible ring homomorphisms, we now need to determine when compatible homomorphisms can be found. In fact, for any *ring* homomorphism ϕ between neural rings it is possible to find at least one compatible homomorphism. Moreover, any compatible homomorphism will act as an extension of the map ϕ to a broader class of functions, as the following theorem states.

Theorem 1.3. *Suppose $\mathcal{C}$ and $\mathcal{D}$ are neural codes on n and m neurons, respectively.*

(i) If $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ is a ring homomorphism, then there exists a ring homomorphism $\tau : R[m] \to R[n]$ which is compatible with ϕ, and thus ϕ is an $R[m]$-module homomorphism.

(ii) For any ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$, the set of compatible ring homomorphisms is exactly the set of ring homomorphisms $\tau : R[m] \to R[n]$ which are extensions of ϕ, in the sense for all $f \in R_{\mathcal{D}}$, $(\phi(f))^{-1}(1) \subseteq (\tau(f))^{-1}(1)$.

Theorem 1.3 is proven in Section 2.2. It shows that compatibility is a powerful descriptive tool. Most notably, since the rings $R[n]$ and $R[m]$ are themselves neural rings, any homomorphism between them is associated to a code map via Theorem 1.1. In Corollary 2.6 we use the second part of Theorem 1.3 to prove that that for a homomorphism τ to be compatible with a homomorphism ϕ, the associated code map q_τ between the full codes $\{0,1\}^n$ and $\{0,1\}^m$ must be an extension of the code map q_ϕ between the smaller codes $\mathcal{C}$ and $\mathcal{D}$.

However, Theorem 1.3 also indicates that by using compatibility to widen our possible allowed module homomorphisms, we obtain again the general class of code maps, similarly to when we considered only the pullback map. That is, any code map corresponds to a ring homomorphism by Theorem 1.1 and thus induces a module map with a compatible ring homomorphism. Hence, every possible code map can be described by a module homomorphism, contrary to our goal of obtaining

only a restricted class of homomorphisms. We have, however, made one small improvement over our original ring homomorphisms: we are now able to track the changes in the activity of neuron i by considering the effect of the image $\tau(x_i)$ of the indicator variable. To restrict to on a class of code maps which contains natural operations but not every possible map, we control the activity of neuron i across the map by placing restrictions on the activity of these indicators under the compatible homomorphisms τ.

The most obvious restriction would be to require that every variable be sent to itself: that $\tau(x_i) = x_i$ for all i for which it is possible, and 0 otherwise. However, this notion is quite restrictive. Most notably, permutation is not an allowable map under this restriction, since the activity of neuron i is forced to remain the same. The only maps we can capture are projection onto the first k neurons, inclusion, and adding 0's to the end of codewords (and compositions thereof). We are not permitted even to delete, say, the second neuron of three, because that would involve renaming neuron 3 as the 'new' neuron 2. See Proposition 2.5 for the formal description of the corresponding code maps.

Since this notion is too restrictive to even capture permutation, we will instead consider a slight relaxation of this restriction which will allow a broader class of code maps.

1.5 Neural ring homomorphism

Instead of restricting to maps τ where $\tau(x_i) = x_i$, we make a slight generalization: we consider *linear-monomial* maps: maps where $\tau(x_i) \in \{x_j, 0, 1\}$. Allowing $\tau(x_i) = x_j$ opens the door for permutation maps. There are other possibilities, as we can map multiple variables to the same, or to constants 0 or 1, so we get more maps than we had under our previous, very restrictive notions. We therefore make the following definition:

Definition. Let $\mathcal{C}$ and $\mathcal{D}$ be neural codes on n and m neurons, respectively. A ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ is a *neural ring homomorphism* if there is a compatible linear-monomial ring homomorphism $\tau : R[m] \to R[n]$.

This is our final choice for a 'best' notion of homomorphism between neural rings, which captures all our desired code maps, and a few extra maps which are still meaningful. Our main result is Theorem 1.4, which shows that neural ring homomorphisms correspond to compositions of the following operations between neural codes.

Theorem 1.4. *ϕ is a neural ring homomorphism if and only if q_ϕ is a composition of the following code maps:*

1. *Permutation of labels*
2. *Adding a codeword (inclusion)*
3. *Deleting the last neuron*
4. *Repeating a neuron*
5. *Adding a trivial neuron (always 0 or 1) to the end of each word.*

We then define a *neural ring isomorphism* to be an isomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{D}}$ which has a compatible linear-monomial isomorphism $\tau : R[m] \to R[n]$, and find that neural ring isomorphisms are exactly permutation maps.

Corollary 1.5. *Neural ring isomorphisms correspond exactly to those code maps which permute the labels on neurons.*

The proofs of Theorem 1.4 and Corollary 1.5 are found in Section 3.1. The set of possible maps described in this result is a restricted class of code maps including our three primary examples of permutation, deletion, and inclusion. Although for deletion the list includes only deletion of the last neuron, note that by composing this operation with permutation we can choose to delete any subset, and project onto the remaining neurons. See Section 3.2 for a detailed description of this, as well as a few examples of other natural maps which these compositions produce. Furthermore, the notion of neural ring isomorphism is a more specialized version of general ring isomorphisms, and which exactly captures permutations. We can repeatedly add codewords to include one code into another. Finally, not every possible code map is captured by neural ring homomorphisms; the first example in Section 1.6 gives an example of a code map which does not correspond to a neural ring homomorphism.

1.6 Examples

We here provide a few examples of code maps, the associated ring homomorphisms and module actions, and determination of whether or not these are neural ring homomorphisms or isomorphisms. Throughout, $\mathcal{B}_n = \langle x_i(1-x_i) \mid i \in [n] \rangle$ is the ideal of Boolean relationships.

1. Let $\mathcal{C} = \{00, 11\}$ and $\mathcal{D} = \{111, 000\}$, and suppose the code map q is a bijection with $q(00) = 111$ and $q(11) = 000$. The pullback map $\phi_q : R_\mathcal{D} \to R_\mathcal{C}$ sends ρ_{111} to ρ_{00} and ρ_{000} to ρ_{11}. The module action on $R_\mathcal{D}$ is given by $x_i \cdot \rho_c = c_i \rho_c$ and extending by linearity. In particular, $x_i \cdot \rho_{000} = 0$ for all i and $x_i \cdot \rho_{111} = \rho_{111}$ for all i, and thus for example , $(x_1 + x_2) \cdot \rho_{111} = x_1 \cdot \rho_{111} + x_2 \cdot \rho_{111} = 2\rho_{111} = 0$. Similarly, the module action on $R_\mathcal{C}$ is given by $x_i \cdot \rho_c = c_i \rho_c$ and extending by linearity.

   ϕ_q is a ring homomorphism; in fact, it is an isomorphism as q is a bijection. However, ϕ_q is is not a neural ring homomorphism. Let $\tau : R[3] \to R[2]$ be a compatible ring homomorphism. Then, we have $\phi_q(x_1 \cdot \rho_{111}) = \phi_q(\rho_{111}) = \rho_{00}$, but also $\phi_q(x_1) \cdot \rho_{111}) = \tau(x_1) \cdot \rho_{00}$, so $\tau(x_1)$ cannot be x_1, x_2, or 0. To be a neural ring homomorphism, we require $\tau(x_1)$ to be linear monomial; the only remaining option is $\tau(x_1) = 1$. But then, $\phi_q(x_1 \cdot \rho_{000}) = \phi_q(0) = 0$, but on the other hand $\phi_q(x_1 \cdot \rho_{000}) = \tau(x_1) \cdot \phi_q(\rho_{000}) = 1 \cdot \rho_{11} = \rho_{11}$ and these are not equal. So there is no compatible τ which is also a linear monomial function on the variables, and hence ϕ_q is not a neural ring homomorphism.

2. Let $\mathcal{C} = \{000, 100, 101, 110\}$ and $\mathcal{D} = \{00, 10, 11, 01\}$. Suppose the code map q is given by $q(000) = 00$, $q(100) = q(101) = 10$ and $q(110) = 11$. The pullback map $\phi_q : R_\mathcal{D} \to R_\mathcal{C}$ sends ρ_{00} to ρ_{000}, ρ_{10} to $\rho_{100} + \rho_{101}$, ρ_{11} to ρ_{110} and ρ_{01} to zero. The module action on $R_\mathcal{D}$ is given by $x_i \cdot \rho_c = c_i \rho_c$ and extending by linearity; for example,

$$(x_1 x_2) \cdot (\rho_{10} + \rho_{11}) = x_1 \cdot (x_2 \cdot \rho_{10}) + x_1 \cdot (x_2 \cdot \rho_{11}) = x_1 \cdot 0 + x_1 \cdot \rho_{11} = \rho_{11}$$

   Similarly, the module action on $R_\mathcal{C}$ is given by $x_i \cdot \rho_c = c_i \rho_c$ and extending by linearity.

   ϕ_q is a ring homomorphism, though not a ring isomorphism; its kernel contains ρ_{01} which is nonzero. It is also a neural ring homomorphism. In fact, there is exactly one possible compatible linear-monomial homomorphism $\tau : R[2] \to R[3]$, given by $\tau(x_1) = x_1$ and $\tau(x_2) = x_2$. We also note that this code map is projection onto the first 2 coordinates (or, deleting the the 3rd neuron) composed with adding the codeword 01.

3. Let $\mathcal{C} = \{0000, 0011, 1000, 1011\}$ and $\mathcal{D} = \{000, 100, 010, 110\}$, and let the code map q be given by $q(0000) = 000$, $q(0011) = 100$, $q(1000) = 010$ and $q(1011) = 110$. The pullback map ϕ_q sends ρ_{000} to ρ_{0000}, ρ_{100} to ρ_{0011}, ρ_{010} to ρ_{1000} and ρ_{110} to ρ_{1011}.

   ϕ_q is a ring homomorphism, and as q is a bijection, we see that $R_\mathcal{C}$ and $R_\mathcal{D}$ are actually isomorphic as rings. ϕ_q is also a neural ring homomorphism; in fact, there are several compatible linear-monomial functions $\tau : R[3] \to R[4]$. One possibility is τ_1 where $\tau_1(x_1) = x_3$, $\tau_1(x_2) = x_1$ and $\tau_1(x_3) = x_2$; this corresponds to the idea that q could be described as deleting neuron 4, and then permuting neurons 1, 2, and 3. Alternatively, we could use τ_2 where $\tau_2(x_1) = x_4, \tau_2(x_2) = x_1$, and $\tau_2(x_3) = 0$; this corresponds to the idea that q could be described as permuting the neurons to the order 4,1,2,3; deleting the last two neurons (2 and 3), and adding a trivial neuron which is always 0. However, while $R_\mathcal{C}$ and $R_\mathcal{D}$ are isomorphic and ϕ_q is a neural ring homomorphism, there is no compatible linear-monomial map τ which is itself an isomorphism (as $R[3]$ and $R[4]$ are not isomorphic) so ϕ_q is not a neural ring isomorphism.

Finally, now that we have seen how examples work at the level of the rings and the code maps, we consider the practical effect of the basic maps given in Theorem 1.4 on the canonical form, one of the major objects of study from our previous work [4].

1.7 Effect on the canonical form

In our original paper [4] we focused our attention on structural information about the code which could be obtained from the ideal $I_\mathcal{C}$, rather than the neural ring $R_\mathcal{C}$. In particular, we split the ideal $I_\mathcal{C}$ into two pieces: $I_\mathcal{C} = \mathcal{B} + J_\mathcal{C}$, where $\mathcal{B}$ covered the standard Boolean relations $\mathcal{B} = \langle x_i^2 - x_i \mid i = 1, ..., n\rangle$. $J_\mathcal{C}$, on the other hand, covered the relationships unique to the code: $J_\mathcal{C} = \langle \rho_v \mid v \in \{0,1\}^n \backslash \mathcal{C}\rangle$. The polynomials which generate $J_\mathcal{C}$ are examples of *pseudo-monomials*: polynomials of the form $\prod_{i\in\sigma} x_i \prod_{j\in\tau}(1 - x_j)$ for $\sigma \cap \tau = \emptyset$. A pseudo-monomial $f \in J_\mathcal{C}$ is *minimal* if there is no other pseudo monomial $g \in J_\mathcal{C}$ of smaller degree so $f = gh$ for some $h \in \mathbb{F}_2[x_1, ..., x_n]$. If we take the set of all minimal pseudo-monomials, we obtain the *canonical form* of $J_\mathcal{C}$, written:

$$CF(J_\mathcal{C}) = \{f \in \mathbb{F}_2[x_1, ..., x_n] \mid f \text{ is a minimal pseudo-monomial of } J_\mathcal{C}\}$$

The canonical form provides not only a condensed set of generators for $J_\mathcal{C}$, but also a minimal combinatorial description of the code itself. It is thus an important and useful object, and so we examine the effect of our basic natural code maps on the canonical form. It's important to note that the effects we examine here, unlike the homomorphisms above, are not contravariant. That is, for a code map $q : \mathcal{C} \to \mathcal{D}$, we describe at how to change $CF(J_\mathcal{C})$ to obtain $CF(J_\mathcal{D})$. The basic code maps have the following effects:

1. Permutation: permuting the neurons requires that we change the canonical form via the same permutation. For example, if neuron 1 is relabeled as neuron 3, then all relationships involving x_1 must now be changed to involve x_3 instead. That is, we perform the same permutation on the labels of the variables.

   Example. Suppose $\mathcal{C} = \{000, 010, 001, 110, 011\}$. Then $CF(J_\mathcal{C}) = \{x_1(1 - x_2), x_1x_3\}$. If we permute the neurons under the permutation (123), then the new code is $\mathcal{D} = \{000, 100, 001, 011, 101\}$, and the new canonical form is $CF(J_\mathcal{D}) = \{x_2(1 - x_3), x_1x_2\}$.

2. Adding codewords: this is the most difficult operation to describe, because the effect of this map can vary greatly depending on the relationship of the codewords being added to the codewords already present. We have, however, written an algorithm which adapts the canonical form iteratively by codeword. In short, all elements of the canonical form which still vanish on the new code are preserved, whereas those that do not are multiplied by a linear term and included if not redundant. See Section 4 for a more detailed description of this algorithm.

   Example. Let $\mathcal{C} = \{000, 010, 001, 110, 011\}$, so $CF(J_\mathcal{C}) = \{x_1(1-x_2), x_1x_3\}$. Suppose we add the codeword 111, so $\mathcal{D} = \{000, 010, 001, 110, 011, 111\}$. Note that $x_1(1-x_2)$ still vanishes on $\mathcal{D}$, but x_1x_3 no longer does, so we must fix it. The 3 linear terms that vanish on 111 are $(1-x_1), (1-x_2)$, and $(1-x_3)$; multiplying x_1x_3 by either $(1-x_1)$ or $(1-x_3)$ won't result in a pseudomonomial, so the only fix that works is $x_1x_3(1-x_2)$. This, however, is redundant, since it is a multiple of $x_1(1-x_2)$. So our new canonical form is simply $CF(J_D) = \{x_1(1-x_2)\}$.

3. Deleting the last neuron: To drop neuron n from the end of each word, we change the canonical form by merely removing all elements of $CF(J_\mathcal{C})$ which involve variable x_n in any way.

   Example. Let $\mathcal{C} = \{000, 010, 001, 110, 011\}$, so $CF(J_\mathcal{C}) = \{x_1(1-x_2), x_1x_3\}$. Suppose we drop neuron 3 from the end of each codeword, and project onto the other two neurons. Then $\mathcal{D} = \{00, 01, 11\}$, and its associated canonical form is $CF(J_\mathcal{D}) = \{x_1(1-x_2)\}$.

4. Repeating a neuron: to repeat the behavior of neuron i with a new neuron x_{n+1}, we duplicate all elements of $CF(J_\mathcal{C})$ which involve x_i, replacing the variable x_i with x_{n+1}. Except in the rare instance where neuron i is trivial (i.e. when either x_i or $(1-x_i)$ is in $J_\mathcal{C}$), we must also add the relations $x_i(1-x_{n+1})$ and $x_{n+1}(1-x_i)$ to the canonical form.

   Example. Let $\mathcal{C} = \{000, 010, 001, 110, 011\}$, so $CF(J_\mathcal{C}) = \{x_1(1-x_2), x_1x_3\}$. Suppose we add a new neuron, neuron 4, which duplicates neuron 2. Then $\mathcal{D} = \{0000, 0101, 0010, 1101, 0111\}$, and its associated canonical form is $CF(J_\mathcal{D}) = \{x_1(1-x_2), x_1(1-x_4), x_1x_3, x_2(1-x_4), x_4(1-x_2)\}$.

5. Adding a trivial neuron: to add a trivial $n+1$st neuron which is always 0 (or always 1) we add the element x_{n+1} (respectively, $1-x_{n+1}$) to the canonical form.

   Example. Let $\mathcal{C} = \{000, 010, 001, 110, 011\}$, so $CF(J_\mathcal{C}) = \{x_1(1-x_2), x_1x_3\}$. Suppose we add a new neuron, neuron 4, which is always 1. Then $\mathcal{D} = \{0001, 0101, 0011, 1101, 0111\}$, and its associated canonical form is $CF(J_\mathcal{D}) = \{x_1(1-x_2), x_1x_3, (1-x_4)\}$.

Knowing how these maps affect the canonical form allows us to transform structural information about the original code quickly and algorithmically into structural information about the resulting image code. Primarily useful here is the algorithm for changing the canonical form in the event of adding a codeword, because it gives us an iterative algorithm for finding the canonical form by building up the code one word at a time, a substantial improvement in efficiency and transparency over our original algorithm from [4]. This algorithm and a proof of its efficacy are found in Section 4.

2 Homomorphisms and code maps

In this section, we give a detailed description of the pullback correspondence between code maps $\mathcal{C} \to \mathcal{D}$ and ring homomorphisms $R_\mathcal{D} \to R_\mathcal{C}$, including the method for determining one from the other.

Recall that the neural ring $R_\mathcal{C}$ for a given code $\mathcal{C}$ is exactly the ring of functions $f : \mathcal{C} \to \{0, 1\}$, and as such is an $\mathbb{F}_2$-vector space. Elements of neural rings may be denoted in different ways: they can be written as polynomials, where it is understood that the polynomial is a representative of its equivalence class mod $I_\mathcal{C}$. Alternatively, using the vector space structure, an element of $R_\mathcal{C}$ can be written as a function which, since each function maps into $\mathbb{F}_2$, is defined completely by the codewords which support it. We will make use of the latter idea frequently, and find it helpful to identify a canonical basis of characteristic functions $\{\rho_c \,|\, c \in \mathcal{C}\}$, where

$$\rho_c(v) = \begin{cases} 1 & v = c \\ 0 & \text{else} \end{cases}$$

In polynomial notation, ρ_c is represented by the polynomial $\prod_{c_i=1} x_i \prod_{c_j=0}(1-x_j)$. The characteristic functions ρ_c form a basis for $R_\mathcal{C}$ as an $\mathbb{F}_2$-vector space, and they have several useful properties; those most important for our proofs are described briefly here.

1. Each element of $R_\mathcal{C}$ corresponds to the formal sum of basis elements for the codewords in its support: $f = \sum_{f(c)=1} \rho_c$.
2. The product of two basis elements is 0 unless they are identical: $\rho_c\rho_d = \begin{cases} \rho_c & c = d \\ 0 & \text{else} \end{cases}$.
3. If $1_\mathcal{C}$ is the identity of $R_\mathcal{C}$, then $1_\mathcal{C} = \sum_{c\in\mathcal{C}} \rho_c$.

These basis elements prove essential to making our proofs intuitive. Recall that given a code map $q : \mathcal{C} \to \mathcal{D}$ the pullback map $q^* : R_\mathcal{D} \to R_\mathcal{C}$ is given by $q^*(f) = f \circ q$; we now show briefly that q^* is a ring homomorphism.

Lemma 2.1. *For any code map $q : \mathcal{C} \to \mathcal{D}$, the pullback $q^* : R_\mathcal{D} \to R_\mathcal{C}$ is a ring homomorphism.*

Proof. First, note that $q^*(f) = f \circ q$ is a well-defined function $\mathcal{C} \to \{0, 1\}$ and any such map is an element of $R_\mathcal{C}$, the ring of all such functions. Now we need to prove $q^* : R_\mathcal{D} \to R_\mathcal{C}$ is a ring homomorphism. To show equality in $R_\mathcal{C}$, we need only show that two functions evaluate the same on any codeword in $\mathcal{C}$. Since elements of $R_\mathcal{C}$ and $R_\mathcal{D}$ are polynomial mappings, we see that both addition and multiplication are preserved:

$$q^*(f + g)(c) = (f + g)(q(c)) = f(q(c)) + g(q(c)) = q^*f(c) + q^*g(c)$$

$$q^*(fg)(c) = (fg)(q(c)) = f(q(c))g(q(c)) = q^*(f)(c)q^*(g)(c).$$

Thus q^* is a ring homomorphism. □

This immediately raises the question: which ring homomorphisms can be obtained as the pullback of code maps? First we make the following observation:

Lemma 2.2. *For any ring homomorphism $\phi : R_\mathcal{D} \to R_\mathcal{C}$, and any element $c \in \mathcal{C}$, there is a unique $d \in \mathcal{D}$ such that $\phi(\rho_d)(c) = 1$.*

Proof. To prove existence, note that $\sum_{c\in\mathcal{C}} \rho_c = 1_{\mathcal{C}} = \phi(1_{\mathcal{D}}) = \phi(\sum_{d\in\mathcal{D}} \rho_d) = \sum_{d\in\mathcal{D}} \phi(\rho_d)$. For each $c \in \mathcal{C}$, $1 = \rho_c(c) = \left(\sum_{c'\in\mathcal{C}} \rho_{c'}\right)(c) = \left(\sum_{d\in\mathcal{D}} \phi(\rho_d)\right)(c)$, and thus $\phi(\rho_d)(c) = 1$ for at least one $d \in \mathcal{D}$. To prove uniqueness, suppose there exist distinct $d, d' \in \mathcal{D}$ such that $\phi(\rho_d)(c) = \phi(\rho_{d'})(c) = 1$. Then we would have $1 = (\phi(\rho_d)\phi(\rho_{d'}))(c) = \phi(\rho_d\rho_{d'})(c) = \phi(0)(c) = 0$, but this is a contradiction. Thus such a d must be unique. □

This allows us to define a unique code map corresponding to any ring homomorphism.

Definition. Given a ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$, we define the associated code map $q_\phi : \mathcal{C} \to \mathcal{D}$ as follows:

$$q_\phi(c) = d_c$$

where d_c is the unique element of $\mathcal{D}$ such that $\phi(\rho_{d_c})(c) = 1$.

Example. Let $\mathcal{C} = \{110, 111, 010, 001\}$ and $\mathcal{D} = \{00, 10, 11\}$. Let $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ be defined by $\phi(\rho_{11}) = \rho_{110} + \rho_{111} + \rho_{010}$, $\phi(\rho_{00}) = \rho_{001}$, and $\phi(\rho_{10}) = 0$. Then the corresponding code map q_ϕ will have $q_\phi(110) = q_\phi(111) = q_\phi(010) = 11$, and $q_\phi(001) = 00$. Note that there is no element $c \in \mathcal{C}$ with $q_\phi(c) = \rho_{10}$ so q_ϕ is not surjective.

We can now prove Theorem 1.1: that code maps and ring homomorphisms are in bijection via the pullback map.

This is a special case of a proposition from [6]. We include our own proof for completeness, and also because it shows constructively how to obtain q_ϕ from ϕ.

Proof of Theorem 1.1. Lemma 2.1 shows that the pullback q^* is a ring homomorphism, so we now prove that any homomorphism can be obtained as the pullback of a code map. Given a ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$, define q_ϕ as above. We must show that the $q_\phi^* = \phi$, and moreover that q_ϕ is the only code map with this property.

The fact that $q_\phi^* = \phi$ holds essentially by construction: let $f \in R_{\mathcal{D}}$, so $f = \sum_{f(d)=1} \rho_d$. Then, for any $c \in \mathcal{C}$,

$$q_\phi^*(f)(c) = f(q_\phi(c)) = \sum_{f(d)=1} \rho_d(q_\phi(c)) = \sum_{f(d)=1} \rho_d(d_c) = \begin{cases} 1 & \text{if } f(d_c) = 1 \\ 0 & \text{if } f(d_c) = 0 \end{cases}$$

whereas, remembering from above that there is exactly one $d \in \mathcal{D}$ such that $\phi(\rho_d)(c) = 1$ and that this d may or may not be in the support of f, we have

$$\phi(f)(c) = \sum_{f(d)=1} \phi(\rho_d)(c) = \begin{cases} 1 & \text{if } d_c \in f^{-1}(1) \\ 0 & \text{if } d_c \notin f^{-1}(1) \end{cases} = \begin{cases} 1 & \text{if } f(d_c) = 1 \\ 0 & \text{if } f(d_c) = 0 \end{cases}.$$

Thus, $\phi = q_\phi^*$.

Finally, to see that q_ϕ is the only code map with this property, suppose we have a different map $q \neq q_\phi$. Then there is some $c \in \mathcal{C}$ with $q(c) \neq q_\phi(c)$; let $d_c = q_\phi(c)$, so $q(c) \neq d_c$. Then $\phi(\rho_{d_c})(c) = 1$ by definition, but $q^*(\rho_{d_c})(c) = \rho_{d_c}(q(c)) = 0$ as $q(c) \neq d_c$. So q^* does not agree with ϕ and hence ϕ is not the pullback of q, so q_ϕ is the unique code map with pullback ϕ. □

While Theorem 1.1 gives a clean bijective correspondence between code maps and ring homomorphisms, it is too general to capture only a class of natural code maps. As the theorem shows, any well-defined code map has a corresponding neural ring homomorphism, even a random assignment $q : \mathcal{C} \to \mathcal{D}$. Our question from this point, therefore, is the following:

Question: What algebraic constraints can be put on homomorphisms between neural rings to capture some meaningful restricted class of code maps?

Since we have determined that ring homomorphism alone is insufficiently restrictive, the next natural idea is to examine ring isomorphism which should be more restrictive. This is true, but the standard notion of isomorphism in fact captures very little actual similarity between codes beyond the number of codewords, as the following lemma shows.

Lemma 2.3. *A ring homomorphism $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ is an isomorphism if and only if $q_\phi : \mathcal{C} \to \mathcal{D}$ is a bijection.*

Proof of Lemma 2.3. Note that $R_{\mathcal{C}} \cong \mathbb{F}_2^{|\mathcal{C}|}$ and $R_{\mathcal{D}} \cong \mathbb{F}_2^{|\mathcal{D}|}$, and $\mathbb{F}_2^{|\mathcal{C}|} \cong \mathbb{F}_2^{|\mathcal{D}|}$ if and only if $|\mathcal{C}| = |\mathcal{D}|$. Suppose ϕ is an isomorphism; then we must have $|\mathcal{C}| = |\mathcal{D}|$. If q_ϕ is not injective, then there is some $d \in \mathcal{D}$ such that $\phi(\rho_d)(c) = 0$ for all $c \in \mathcal{C}$. But then $\phi(\rho_d) = 0$, which is a contradiction since ϕ is an isomorphism so $\phi^{-1}(0) = \{0\}$. Thus q_ϕ is injective, and since $|\mathcal{C}| = |\mathcal{D}|$, this means q_ϕ is a bijection.

On the other hand, suppose $q_\phi : \mathcal{C} \to \mathcal{D}$ is a bijection. Then $|\mathcal{C}| = |\mathcal{D}|$, so $R_{\mathcal{C}} \cong R_{\mathcal{D}}$, and as both are finite, $|R_{\mathcal{C}}| = |R_{\mathcal{D}}|$. Consider an arbitrary element $f \in R_{\mathcal{C}}$. For each $c \in f^{-1}(1)$, there is a unique $d \in \mathcal{D}$ so $\phi(\rho_d) = c$; furthermore as q_ϕ is a bijection, all these d are distinct. . Then

$$\phi\Big(\sum_{\substack{d=q_\phi(c),\\ c\in f^{-1}(1)}} \rho_d\Big) = \sum_{\substack{d=q_\phi(c)\\ c\in f^{-1}(1)}} \phi(\rho_d) = \sum_{c\in f^{-1}(1)} \rho_c = f.$$

Hence ϕ is surjective, and since $|R_{\mathcal{C}}| = |R_{\mathcal{D}}|$, we know ϕ is also bijective and hence an isomorphism. □

This shows us that not only would two permutation-equivalent codes have isomorphic neural rings, but in fact *any* two codes with the same number of codewords have isomorphic neural rings. Lemma 2.3 highlights the main difficulty with using ring homomorphism and isomorphism alone: the neural rings are rings of functions from $\mathcal{C}$ to $\{0, 1\}$, and the abstract structure of such a ring depends solely upon the size of the code, $|\mathcal{C}|$. By considering such rings abstractly without their presentation, we can reflect no other structure, not even such basic information as the number of bits per codeword. In particular, we cannot track the behavior of the variables x_i which represent individual neurons, since once we move to an abstract view of the ring the presentation is lost, and it is merely an $\mathbb{F}_2$-vector space determined completely by the number of codewords.

In order to track the activity of neurons, we turn next to modules, where we can consider the neural ring *with* its presentation, tracking the variables which represent the neurons via their action on elements of the neural ring. We will consider each neural ring as a carefully designed module under the ring for the full code on the same number of neurons, which will allow us to keep track of the structure.

2.1 Neural rings as modules

Given a code $\mathcal{C}$ on n neurons, we consider $R_{\mathcal{C}}$ as a module under $R[n]$ under the module action

$$R[n] \times R_{\mathcal{C}} \to R_{\mathcal{C}}$$

given by multiplication - that is, for a codeword $c \in \mathcal{C}$, we have

$$(r \cdot f)(c) = r(c)f(c).$$

In particular, for any basis element ρ_c, we have

$$x_i \cdot \rho_c = \begin{cases} \rho_c & \text{if } c_i = 1 \\ 0 & \text{if } c_i = 0 \end{cases}.$$

and we can determine the rest of the action by extension. In particular, for any $r \in R[n]$ and $f \in R_{\mathcal{C}}$, we will have $(r \cdot f)^{-1}(1) = (r^{-1}(1) \cap f^{-1}(1))$.

This module action is equivalent to multiplication of polynomials, with the result considered as an element of $R_{\mathcal{C}}$. In particular, we always impose the relationships $x_i(1 - x_i) = 0$, which also means $x_i^2 = x_i$ and $(1 - x_i)^2 = (1 - x_i)$. The use of these relationships can be seen more clearly in the following example:

Example. Consider again the code $\mathcal{C} = \{000, 100, 101\}$. $R_{\mathcal{C}}$ is a module under $R[3] = R_{\{0,1\}^3}$.

- Consider the element $1 - x_1$ of $R[3]$. Then

  $$(1 - x_1) \cdot (1 - x_1)(1 - x_2)(1 - x_3) = (1 - x_1)(1 - x_2)(1 - x_3).$$

  whereas

  $$(1 - x_1) \cdot x_1(1 - x_2)(1 - x_3) = 0.$$

- Consider the element x_1x_2 of $R[3]$. Although x_1x_2 is a nonzero element in $R[3]$, it evaluates to 0 for all codewords in $\mathcal{C}$, so for any element $f \in R_{\mathcal{C}}$, we have $x_1x_2 \cdot f = 0$.

As another way to look at this action, we can consider each neural ring using the basis of characteristic functions and perform the same calculations. $R[3]$ is itself a neural ring, and so has basis elements $\{\rho_c \,|\, c \in \{0,1\}^3\}$. We can look at the action in terms of these basis elements (using the same examples as above)

- The element $1 - x_1 \in R[3]$ is written $\rho_{000} + \rho_{010} + \rho_{001} + \rho_{011}$. In $R_{\mathcal{C}}$, we have $(1 - x_1)(1 - x_2)(1 - x_3) = \rho_{000}$ and $x_1(1 - x_2)(1 - x_3) = \rho_{100}$. Then

  $$(\rho_{000} + \rho_{010} + \rho_{001} + \rho_{011}) \cdot \rho_{000} = \rho_{000}$$

  whereas

  $$(\rho_{000} + \rho_{010} + \rho_{001} + \rho_{011}) \cdot \rho_{011} = 0.$$

- In $R[3]$, $x_1x_2 = \rho_{110} + \rho_{111}$. Although this is a nonzero element in $R[3]$, 110 and 111 are not in $\mathcal{C}$ and thus $R_{\mathcal{C}}$ has no basis elements for them. Thus, for any element $f \in R_{\mathcal{C}}$, we have $(\rho_{110} + \rho_{111}) \cdot f = 0$.

We have already observed that a neural ring is an $\mathbb{F}_2$-vector space with basis $\{\rho_c \,|\, c \in \mathcal{C}\}$. As an $R[n]$-module, the basis elements are the unique elements of $R_{\mathcal{C}}$ (besides 0) such that $x_i \cdot f \in \{f, 0\}$ for all $i \in [n]$. We also see that any two basis elements can be distinguished by the action of some x_i; that is, for $\rho_c \neq \rho_{c'}$, there is some i so $c_i \neq c'_i$, and hence exactly one of $x_i \cdot \rho_c$ and $x_i \cdot \rho_{c'}$ will be 0. In fact, these characteristics are sufficient to characterize the neural rings as $R[n]$-modules. This is the substance of Theorem 1.2, which we now prove.

Proof of Theorem 1.2. Let M be an $R[n]$-module with $\mathbb{F}_2$-basis $\rho_1, ..., \rho_d$ with the given properties. First, we identify the basis elements ρ_j, if not already known. We can do this by considering the action $x_i \cdot m$ on every element of M. Note that if m is neither a basis element ρ_j, nor 0_M, then

it is a formal sum of at least two basis elements; by property 2, we can distinguish between these two elements with some x_i, so $x_i \cdot m$ will be neither m nor 0. Hence, the set $\{m \in M \mid x_i \cdot m \in \{m, 0\}$ for all $i \in [n], m \neq 0_M\}$ will be exactly the set $\{\rho_1, ..., \rho_d\}$. It is possible that a basis element ρ_j may be such that $x_i \cdot \rho_j = 0$ for all $i \in [n]$, but this basis element is not the same as the identity 0_M and we distinguish them if necessary by considering the action of $1 - x_i$.

Once we have the set of basis elements $\rho_1, ..., \rho_d$, we associate a related code $\mathcal{C} \subset \{0,1\}^n$. To each basis element ρ_j we associate a unique word c^j by taking $c^j_i = \begin{cases} 1 & x_i \cdot \rho_j = \rho_j \\ 0 & x_i \cdot \rho_j = 0 \end{cases}$ This assignment is well-defined by property 1, and unique by property 2. We then take $\mathcal{C} = \{c^j \mid j \in [d]\}$. Then we make the obvious choice of module isomorphism: $\phi : M \to R_{\mathcal{C}}$ given by $\phi(\rho_j) = \rho_{c^j}$. By our choice of associated codewords, we have ensured that the module action is preserved across ϕ: that is, $\phi(x_i \cdot \rho_j) = x_i \cdot \rho_{c^j}$, and as both modules are $\mathbb{F}_2$-vector spaces, they are isomorphic as $R[n]$ modules.

On the other hand, suppose M is an $R[n]$-module which is isomorphic to a neural ring $R_{\mathcal{C}}$; let $\phi : R[n] \to M$ be a module isomorphism. Since $R_{\mathcal{C}}$ is a finite $\mathbb{F}_2$-vector space, we know M is an $\mathbb{F}_2$-vector space with basis $\{\phi(\rho_c) \mid c \in \mathcal{C}\}$; the relevant properties of the basis follow immediately by properties of module isomorphism. □

Example. Suppose M is an $R[3]$-module with 3 basis elements ρ_1, ρ_2, ρ_3. Suppose, for example, that $x_1 \cdot \rho_1 = \rho_1$ and $x_2 \cdot \rho_1 = x_3 \cdot \rho_1 = 0$. We thus know ρ_1 corresponds to the word 100 and hence $\rho_1 \leftrightarrow \rho_{100}$. Similarly, suppose we then find that $\rho_2 \leftrightarrow \rho_{000}$ and $\rho_3 \leftrightarrow \rho_{101}$; then $M \cong R_{\mathcal{C}}$ as $R[n]$-modules, where $\mathcal{C} = \{000, 100, 101\}$.

Now that we have a framework to consider a neural ring $R_{\mathcal{C}}$ on n neurons as an $R[n]$-module that preserves the code structure, we consider module homomorphisms between neural rings as an alternative to the too-general ring homomorphisms

Definition. Let R be a ring and suppose M, N are R-modules. Then a map $\phi : M \to N$ is an *R-module homomorphism* if for any $m \in M, n \in N$, and $r, s \in R$, we have

$$f(r \cdot m + s \cdot n) = r \cdot f(m) + s \cdot f(n).$$

A module homomorphism ϕ is an *R-module isomorphism* if ϕ is a bijection.

For neural rings, this is an incredibly restrictive notion, because we can only consider module homomorphisms between $R_{\mathcal{C}}$ and $R_{\mathcal{D}}$ when the codes $\mathcal{C}$ and $\mathcal{D}$ have the same length n and are thus both $R[n]$-modules. Under this framework, the code maps we are allowed are only inclusion.

Lemma 2.4. *Given a code map $q : \mathcal{C} \to \mathcal{D}$, where $\mathcal{C}, \mathcal{D}$ are both subsets of $\{0,1\}^n$, the induced ring homomorphism ϕ_q is a module homomorphism $R_{\mathcal{D}} \to R_{\mathcal{C}}$ if and only if q is an inclusion map $q : \mathcal{C} \hookrightarrow \mathcal{D}$, and it is a module isomorphism if and only if q is the identity, i.e. $\mathcal{C} = \mathcal{D}$ and $q(c) = c$ for all c.*

Proof of Lemma 2.4. Because we have an $\mathbb{F}_2$-vector space structure, ϕ_q is a module homomorphism if and only if for every $i \in [n]$ and $d \in \mathcal{D}$, we have

$$x_i \cdot \phi_q(\rho_d) = \phi_q(x_i \cdot \rho_d).$$

These are equal if and only if they evaluate the same on every element $c \in \mathcal{C}$. Computing, we see

$$x_i \cdot \phi_q(\rho_d)(c) = \begin{cases} 1 & \text{if } c_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } c_i = 0 \text{ or } q(c) \neq d \end{cases}$$

whereas

$$\phi_q(x_i \cdot \rho_d)(c) = \begin{cases} 1 & \text{if } d_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } d_i = 0 \text{ or } q(c) \neq d \end{cases}$$

and these are equal on every $c \in \mathcal{C}$ if and only if whenever $q(c) = d$, we have $c_i = d_i$ for all $i \in [n]$. This is only true if q is an inclusion map.

Finally, we have an isomorphism if and only if ϕ_q is a bijection, but this happens if and only if q is itself a bijection, as we saw with ring homomorphisms. Thus ϕ_q is a module isomorphism if and only if q is the identity map. □

Clearly, this notion is far too restrictive; in particular, code maps between two codes of different lengths are not permitted at all; nor are permutations possible. In the next section, we show how to extend our list of possible maps further by allowing module homomorphisms between $R_{\mathcal{D}}$ and $R_{\mathcal{C}}$ under *different* rings $R[n]$ and $R[m]$.

2.2 Modules under different rings

As we discussed in Section 1.3, we may consider an $R[n]$-module $R_{\mathcal{C}}$ as an $R[m]$ module, provided we have a ring homomorphism $\tau : R[m] \to R[n]$.

For any given n, m there is a standard map $\tau : R[m] \to R[n]$:

$$\tau(x_i) = \begin{cases} x_i & \text{if } i \leq n \\ 0 & \text{if } i > n \end{cases}$$

Now, our question becomes: which code maps induce $R[m]$-module homomorphisms under this construction? Unfortunately, we still find the class of allowable code maps to be very restrictive, and dependent almost entirely on the relative sizes of m and n:

Proposition 2.5. *A code map $q : \mathcal{C} \to \mathcal{D}$ where $\mathcal{C} \subset \{0,1\}^n$ and $\mathcal{D} \subset \{0,1\}^m$ induces an $R[m]$-module homomorphism $\phi_q : R_{\mathcal{D}} \to R_{\mathcal{C}}$ under the map τ described above if and only if*

1. *$[n > m]$ q is the deletion of the last $n - m$ neurons, possibly composed with an inclusion map.*
2. *$[n = m]$ q is an inclusion map.*
3. *$[n < m]$ q is given by adding $m - n$ trivial neurons (never firing, always 0) to the end each word, possibly composed with an inclusion map.*

Proof of Proposition 2.5. All parts proceed very like the proof of Lemma 2.4.

1. ϕ_q is a module homomorphism if and only if for every $i \in [m]$ and $d \in \mathcal{D}$, we have

$$\tau(x_i) \cdot \phi_q(\rho_d) = \phi_q(x_i \cdot \rho_d).$$

These are equal if and only if they evaluate the same on every element $c \in \mathcal{C}$. For all $i \in [m]$,

$$\tau(x_i) \cdot \phi_q(\rho_d)(c) = x_i \cdot \phi_q(\rho_d)(c) = \begin{cases} 1 & \text{if } c_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } c_i = 0 \text{ or } q(c) \neq d \end{cases}$$

whereas

$$\phi_q(x_i \cdot \rho_d)(c) = \begin{cases} 1 & \text{if } d_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } d_i = 0 \text{ or } q(c) \neq d \end{cases}$$

and these are equal on every $c \in \mathcal{C}$ if and only if whenever $q(c) = d$, we have $c_i = d_i$ for all $i \in [m]$. This exactly determines the map given by deleting the last $n - m$ neurons, as $\mathcal{C} \subset \{0,1\}^n$, $\mathcal{D} \subset \{0,1\}^m$, and c agrees with $q(c)$ on the first m neurons. However, since we place no restriction on ϕ_q being a bijection, inclusion may also be permitted; i.e., it is possible that $q(\mathcal{C}) \subsetneq \mathcal{D}$.

2. If $n = m$, then we are exactly in the case of Lemma 2.4 so the only possible map is the inclusion map $q : \mathcal{C} \hookrightarrow \mathcal{D}$.

3. ϕ_q is a module homomorphism if and only if for every $i \in [m]$ and $d \in \mathcal{D}$, we have

$$\tau(x_i) \cdot \phi_q(\rho_d) = \phi_q(x_i \cdot \rho_d).$$

These are equal if and only if they evaluate the same on every element $c \in \mathcal{C}$. Computing for $i \leq n$, we have

$$\tau(x_i) \cdot \phi_q(\rho_d)(c) = x_i \cdot \phi - q(\rho_d)(c) = \begin{cases} 1 & \text{if } c_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } c_i = 0 \text{ or } q(c) \neq d \end{cases}$$

whereas

$$\phi_q(x_i \cdot \rho_d)(c) = \begin{cases} 1 & \text{if } d_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } d_i = 0 \text{ or } q(c) \neq d \end{cases}.$$

and these are equal on every $c \in \mathcal{C}$ and if and only if whenever $q(c) = d$, we have $c_i = d_i$ for all $i \in [n]$.

For $i > n$, we know

$$\tau(x_i) \cdot \phi_q(\rho_d)(c) = 0 \cdot \phi_q(\rho_d)(c) = 0$$

whereas

$$\phi_q(x_i \cdot \rho_d)(c) = \begin{cases} 1 & \text{if } d_i = 1 \text{ and } q(c) = d \\ 0 & \text{if } d_i = 0 \text{ or } q(c) \neq d \end{cases}.$$

so these are equal on every $c \in \mathcal{C}$ if and only if whenever $q(c) = d$ we have $d_i = 0$ for all $i > n$.

These two requirements combine to describe the map given by adding $m - n$ zeros to the end of each codeword $c \in \mathcal{C}$. As in (1), we also must allow inclusion since we are not requiring that ϕ_q to be a bijection and it is possible that $q(\mathcal{C}) \subsetneq \mathcal{D}$.

Hence, these are the only 3 possible types of maps under the canonical map τ. □

Once again, this set of allowable code maps is far too restrictive; in particular, permutation maps are still omitted. However, we have obtained the ability to consider module homomorphisms between neural rings of different length, so this construction is a good step towards our goal. We therefore consider situations where this map τ is allowed to be more general, not just the canonical map described above.

As a code map $q : \mathcal{C} \to \mathcal{D}$ determines the related map ϕ_q between the modules $R_\mathcal{D}$ and $R_\mathcal{C}$, the code map q places certain restrictions on which ring homomorphisms τ may be used between $R[m]$ and $R[n]$. Since we have a preconceived idea of which code maps are desirable (permutation, dropping a neuron, inclusion), it is natural to let the code maps drive the choice of ring homomorphism. The ring homomorphisms for which a certain map ϕ_q is still a module homomorphism are the *compatible* homomorphisms, which we formally define in Section 1.4.

It is worth noting that not every possible function between two neural rings has a compatible ring homomorphism.

Example. Consider the codes $\mathcal{C} = \{000, 100, 101\}$ and $\mathcal{D} = \{00, 10, 11\}$, and let the map $\phi : R_\mathcal{D} \to R_\mathcal{C}$ be given by $\phi(\rho_{00}) = \rho_{000} + \rho_{100}$, $\phi(\rho_{10}) = \rho_{100}$, and $\phi(\rho_{11}) = 0$. Extending by linearity to all elements of $R_\mathcal{D}$ gives us a group homomorphism, which is easy to check, although not a ring homomorphism (also easy to check).

As a polynomial map, this is the group homomorphism given by: $x_1 \to y_1(1 - y_2)(1 - y_3)$, $x_2 \to 0$. There is, however, no compatible ring homomorphism $\tau : R[2] \to R[3]$ so that ϕ is an $R[2]$-module homomorphism. To see this, note that any such homomorphism τ would require $\phi(\rho_{00} \cdot \rho_{00}) = \tau(\rho_{00})\phi(\rho_{00}) = \tau(\rho_{00}) \cdot [\rho_{000} + \rho_{100}]$. But as $\rho_{00} \cdot \rho_{00} = \rho_{00}$, this must equal $\rho_{000} + \rho_{100}$. So $\tau(\rho_{00})$ must preserve ρ_{000} and ρ_{100}. Similarly, $\tau(\rho_{10})$ must preserve ρ_{100}. So $\tau(\rho_{10})\tau(\rho_{00})$ must preserve ρ_{100} at least, so $\tau(\rho_{10})\tau(\rho_{00}) \neq 0$. Note $\tau(\rho_{00}\rho_{10}) = \tau(0) = 0$, but as τ is a ring homomorphism, we also have $\tau(\rho_{00}\rho_{10}) = \tau(\rho_{00})\tau(\rho_{10}) \neq 0$. So no such τ can exist; there is no compatible ring homomorphism for ϕ.

However, *ring* homomorphisms between two neural rings are guaranteed to have at least one compatible ring homomorphism τ; this is the substance of the first part of Theorem 1.3. For this result, we use the idea that elements of $R_\mathcal{D}$ can also be thought of as elements of the ambient ring $R[m]$. For example, each basis element ρ_d of $R_\mathcal{D}$ is the function which detects only the codeword d; however, since $d \in \{0, 1\}^m$, we know $R[m]$ has a basis element ρ_d which detects only d as well, and we associate these two basis elements. Likewise, any function $f \in R_\mathcal{D}$ corresponds the subset $f^{-1}(1) \subset \mathcal{D}$ which it detects, so we can consider f as a function in $R[m]$ which detects the same set of codewords. Theorem 1.3 not only states that any ring homomorphism ϕ_q between neural rings has at least one compatible homomorphism, but also states that the set of compatible ring homomorphisms is given those maps which agree as far as possible with the original ϕ_q. We now prove Theorem 1.3.

Proof of Theorem 1.3. Let $\phi : R_\mathcal{D} \to R_\mathcal{C}$ be a ring homomorphism. To construct a compatible ring homomorphism $\tau : R[m] \to R[n]$, first select one basis element ρ_d of $R_\mathcal{D}$. Note that ρ_d (as the function which detects exactly the codeword $\{d\}$) is also a basis element of $R[m]$, and define $\tau(\rho_d) = \phi(\rho_d) + \sum_{v \in \{0,1\}^n \setminus \mathcal{C}} \rho_v$; that is, $\tau(\rho_d)$ will detect all the same codewords as $\phi(\rho_d)$, but also all the codewords of $\{0, 1\}^n$ which are not part of $\mathcal{C}$. For all other $d \in \mathcal{D}$, define $\tau(\rho_d) = \phi(\rho_d)$, and for all $v \in \{0, 1\}^m \setminus \mathcal{D}$, define $\tau(\rho_v) = 0$. Extend τ to all elements of $R[m]$ by linearity; that is, if $f = \sum \rho_c$, then $\tau(f) = \sum \tau(\rho_c)$. This gives a ring homomorphism.

Now, we will show that the property of compatibility is equivalent to the property of extensions. Let $\tau : R[m] \to R[n]$ be a ring homomorphism, and let $f = \sum_{d \in f^{-1}(1)} \rho_d \in R_\mathcal{D}$. As $f \cdot f = f$, then τ is compatible with ϕ if and only if we have $\phi(f) = \phi(f \cdot f) = \tau(f) \cdot \phi(f)$, which occurs if and only if $\tau(f)$ detects at least the same codewords as $\phi(f)$, which happens if and only if $\tau(f)^{-1}(1) \supseteq \phi(f)^{-1}(1)$. □

Since our ambient rings $R[n]$ and $R[m]$ are themselves neural rings, we note that that each map $\tau : R[m] \to R[n]$ is also a ring homomorphism between the neural rings for the complete

codes. We have shown that code maps correspond to ring homomorphisms, and thus each ring homomorphism τ corresponds to a unique code map $q_\tau : \{0,1\}^m \to \{0,1\}^n$ between the complete codes. Furthermore, Theorem 1.3 shows that τ is compatible with ϕ if and only if it is an extension of ϕ, in that $\phi(f)^{-1}(1) \subseteq \tau(f)^{-1}(1)$. We also know from Theorem 1.3 that given a ring homomorphism $\phi : R_\mathcal{D} \to R_\mathcal{C}$, we can always find a ring homomorphism $\tau : R[m] \to R[n]$ which is compatible with ϕ. We now confirm that idea by considering the code maps and prove the following Lemma, which shows we can generate a compatible τ at the level of the codes. We do this by taking any code map $q' : \{0,1\}^n \to \{0,1\}^m$ which extends q (so $q'(c) = q(c)$ for all $c \in \mathcal{C}$) and using the corresponding ring homomorphism τ_q. Then $\tau_{q'}$ will take each function to its pullback by q', and it will be compatible with ϕ.

Corollary 2.6. *The ring homomorphism $\tau : R[m] \to R[n]$ is compatible with the ring homomorphism $\phi : R_\mathcal{D} \to R_\mathcal{C}$ if and only if $q_\phi = q_\tau\big|_\mathcal{C}$.*

Proof. Suppose τ is compatible with ϕ. Note that $q_\tau\big|_\mathcal{C} = q_\phi$ if and only if $q_\tau(c) = q_\phi(c)$ for all $c \in \mathcal{C}$. So, suppose by way of contradiction that $q_\tau(c) \neq q_\phi(c)$ for some $c \in \mathcal{C}$. Let $d = q_\phi(c)$. Then, $\phi(\rho_d)(c) = 1$. But $\phi(\rho_d) = \phi(\rho_d \cdot \rho_d) = \tau(\rho_d) \cdot \phi(\rho_d)$, and we know $\tau(\rho_d)(c) = 0$, not 1, so τ and ϕ cannot be compatible. This is a contradiction.

Now, suppose $q_\phi = q_\tau\big|_\mathcal{C}$. Suppose $\phi(\rho_d)(c) = (1)$. By our code map-homomorphism correspondence, this means that $q_\phi(c) = d$. So then $q_\tau(c) = d$ also, and thus again by the correspondence, $\tau(\rho_d)(c)(1)$. Thus, for each $d \in \mathcal{D}$, we have $\tau(\rho_d)^{-1}(1) \supseteq \phi(\rho_d)^{-1}(1)$, and by Theorem 1.3, τ is compatible with ϕ. □

Example. Consider again the codes $\mathcal{C} = \{000, 100, 101\}$ and $\mathcal{D} = \{00, 10, 11\}$. Let the map $\phi : R_\mathcal{D} \to R_\mathcal{C}$ be given by $\phi(\rho_{00}) = \rho_{000}$, $\phi(\rho_{10}) = \rho_{100} + \rho_{101}$, and $\phi(\rho_{11}) = 0$, and extend by linearity to all other elements. It is easy to check that ϕ is a ring homomorphism.

Consider $\tau_1 : R[2] \to R[3]$ given by $\tau_1(\rho_{00}) = \rho_{000} + \rho_{001} + \rho_{010} + \rho_{110} + \rho_{011} + \rho_{111}, \tau_1(\rho_{10}) = \rho_{100} + \rho_{101}$, and $\tau_1(\rho_{11}) = 0$. Extend again by linearity. Then τ_1 is a compatible ring homomorphism, which is again quick to check.

Now, consider $\tau_2 : R[2] \to R[3]$ given by $\tau_2(\rho_{00}) = \rho_{000} + \rho_{001} + \rho_{010} + \rho_{110} + \rho_{011} + \rho_{111}, \tau_2(\rho_{10}) = \rho_{100}$, and $\tau_2(\rho_{11}) = \rho_{101}$. τ_2 is not a compatible ring homomorphism, as if it were, we would have

$$\rho_{100} + \rho_{101} = \phi(\rho_{10}) = \phi(\rho_{10}\rho_{10}) = \tau_2(\rho_{10})\phi(\rho_{10}) = \rho_{100} \cdot (\rho_{100} + \rho_{101}) = \rho_{100}$$

which is a contradiction.

Theorem 1.3, along with Corollary 2.6 show that compatibility alone is too general a notion. While restricting our τ to the standard map $x_i \mapsto x_i$ was far too restrictive, allowing τ to be *any* compatible map once again gives us all possible code maps. In Theorem 1.1 it was established that ring homomorphisms $\phi : R_\mathcal{D} \to R_\mathcal{C}$ between two neural rings are in correspondence with the set of possible functions $q : \mathcal{C} \to \mathcal{D}$. Theorem 1.3 makes it clear that even considering module homomorphisms does not capture the restricted class of code maps we desire, since not every code map preserves structure, but every code map generates a related ring homomorphism and therefore a related module homomorphism. Using module properties we can extract code structure and track its behavior across the homomorphism, but we we have not yet ensured that structure is preserved across maps.

3 Linear-monomial module homomorphisms & neural ring homomorphisms

We have seen from Proposition 2.5 that if we require indeterminates map to themselves if possible, and 0 if not, we obtain a very restricted class of code maps. In the rest of Section 2, we proved that if we make no requirements on the indeterminates and take module homomorphisms under any compatible homomorphism, we can get any possible code map. What we are working towards is something between these ideas - a set of code maps where neurons preserve their basic activity, but we might only look at a subset of the neurons, or we might allow their names to change.

Therefore, we now place make a slight relaxation on the model from Proposition 2.5 and consider *linear monomial* homomorphisms - where $\tau(x_i) \in \{x_j, 0, 1\}$, as discussed in Section 1.5. The code maps which have a possible compatible homomorphism which is also linear-monomial we will choose as our *neural ring homomorphisms*. Not all ring homomorphisms are linear-monomial, as shown by the following example:

Example. Consider the map $\tau : R[1] \to R[2]$ given by $\tau(\rho_1) = \rho_{01} + \rho_{10}$ and $\tau(\rho_0) = \rho_{00} + \rho_{11}$. Here $\tau(x_1) = x_1 + x_2$, and $\tau(x_i) \notin \{x_1, x_2, 0, 1\}$ as would be required.

Observe that a linear-monomial ring homomorphism $\tau : R[m] \to R[n]$ is defined by a vector $S = (s_1, ..., s_m)$, where $S_i \in [n] \cup \{0, u\}$, so that

$$\tau(x_i) = \begin{cases} x_j & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = u \end{cases}.$$

S is a vector which stores the pertinent information about τ, and each possible S with $s_i \in [n] \cup \{0, u\}$ defines a possible neuron-preserving τ. We refer to the τ defined by S as τ_S.

Remark: We have defined linear-monomial homomorphism at the level of maps between the ambient rings $R[n]$ only, and not between neural rings in general, even though the indeterminate x_i does have meaning in a general neural ring. This allows us to have a linear-monomial homomorphism τ for any given S without accidentally contradicting given relationships; that is, we don't need to consider relationships amongst the x_i's in $R[n]$, since there are no relationships in $R[n]$ besides the Boolean relationships $x_i(1 - x_i) = 0$. Thus, any choice of the vector S will give a valid ring homomorphism.

Lemma 3.1. *The composition of two linear-monomial homomorphisms is also linear-monomial.*

Proof of Lemma 3.1. Suppose $S = (s_1, ..., s_n)$ with $s_i \in [m] \cup \{0, u\}$ and $T = (t_1, ..., t_m)$ with $t_i \in [\ell] \cup \{0, u\}$ are given as above, with $\tau_S : R[n] \to R[m]$ and $\tau_T : R[m] \to R[\ell]$. To prove the lemma, we need to find $W = (w_1, ..., w_n)$ with $w_i \in [\ell] \cup \{0, u\}$ so $\tau_W = \tau_T \circ \tau_S : R[n] \to R[\ell]$.

Define the vector W by

$$w_i = \begin{cases} t_{s_i} & \text{if } s_i \in [m] \\ 0 & \text{if } s_i = 0 \\ u & \text{if } s_i = u \end{cases}$$

We use variables z_i for $R[n]$, y_i for $R[m]$, x_i for $R[\ell]$ to make it clear at all times what ring we

are working in. Then, unraveling the definitions,

$$\tau_W(z_i) = \begin{cases} x_j & \text{if } t_{s_i} = j \\ 0 & \text{if } t_{s_i} = 0 \\ & \text{or if } s_i = 0 \\ 1 & \text{if } t_{s_i} = u \\ & \text{or if } s_i = u \end{cases} = \begin{cases} x_j & \text{if } s_i = k \text{ and } t_k = j \\ 0 & \text{if } s_i = k \text{ and } t_k = 0 \\ & \text{or if } s_i = 0 \\ 1 & \text{if } s_i = k \text{ and } t_k = u \\ & \text{or if } s_i = u \end{cases}$$

$$= \begin{cases} x_j & \text{if } \tau_S(z_i) = y_k \text{ and } \tau_T(y_k) = x_j \\ 0 & \text{if } \tau_S(z_i) = y_k \text{ and } \tau_T(y_k) = 0 \\ & \text{or if } \tau_S(z_i) = 0 \\ 1 & \text{if } \tau_S(z_i) = y_k \text{ and } \tau_T(y_k) = 1 \\ & \text{or if } \tau_S(z_i) = 1 \end{cases} = \begin{cases} x_j & \text{if } \tau_T \circ \tau_S(z_i) = x_j \\ 0 & \text{if } \tau_T \circ \tau_S(z_i) = 0 \\ 1 & \text{if } \tau_T \circ \tau_S(z_i) = 1 \end{cases}.$$

□

Given such a vector $S = (s_1, ..., s_m)$, where $s_i \in [n] \cup \{0, u\}$, define the code map $q_S : \{0,1\}^n \to \{0,1\}^m$ to be given by $q_S(c) = d$ if and only if

$$d_i = \begin{cases} c_j & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = u \end{cases}.$$

Lemma 3.2. *Let $\mathcal{C} = \{0,1\}^n$ and $\mathcal{D} = \{0,1\}^m$, and suppose $q_S : \mathcal{C} \to \mathcal{D}$ is defined by S as above. Then $\phi_{q_S} = \tau_S$, and thus $q_S = q_{\tau_S}$ by Corollary 2.6.*

Proof of Lemma 3.2. Write $q_S = q$ and $\tau_S = \tau$. Note $\phi_q(y_i) = q^* y_i$, so for any $c \in \mathcal{C}$,

$$q^* y_i(c) = y_i(q(c)) = d_i = \begin{cases} c_j & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = u \end{cases}.$$

On the other hand, we also have

$$\tau(y_i)(c) = \begin{cases} x_j(c) & \text{if } \tau(y_i) = x_j \\ 0 & \text{if } \tau(y_i) = 0 \\ 1 & \text{if } \tau(y_i) = 1 \end{cases} = \begin{cases} x_j(c) & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = u \end{cases} = \begin{cases} c_j & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = 1 \end{cases}.$$

Thus τ and ϕ_q are identical on $\{y_i\}$, and hence are identical everywhere. □

Here we show how our natural motivating code maps, as well as some other types, are defined by S-vectors. Throughout, let $\mathcal{C}$ be the full code on n neurons and $c = (c_1, ..., c_n)$ an element of $\mathcal{C}$. $\mathcal{D}$ will be the full code on m neurons, but the relationship of m and n will vary for each example.

1. **Permuting the labels**: Let $\sigma \in \mathcal{S}_n$ be a permutation. To relabel the code so neuron i is relabeled $\sigma(i)$, we use $S = (\sigma(1), ..., \sigma(n))$. Then $q_S(c) = d$, where $d = (c_{\sigma(1)}, ..., c_\sigma(n))$. We require $q_S(\mathcal{C}) = \mathcal{D}$.

2. **Adding a codeword**: Let $S = (1, 2, ..., n)$. This defines an inclusion map, so $q(c) = c$. We may use this anytime we have $q(\mathcal{C}) \subseteq \mathcal{D}$. Then all codewords in $\mathcal{D} \setminus q(\mathcal{C})$ are "added." Again this is exactly the same as our inclusion maps defined under the canonical τ with $n = m$.

3. **Deleting the last neuron**: Let $S = (1, 2, ..., n-1)$. Then $q_S(c) = d$, where $d = (c_1, ..., c_{n-1})$. We require $q_S(\mathcal{C}) = \mathcal{D}$. In particular, this is the map given by the canonical τ, where $m = n-1$.

4. **Adding a neuron which repeats neuron** i to the end of each word: $S = (1, 2, ..., n, i)$. Then $q_S(c) = d$, where $d = (c_1, ..., c_n, c_i)$. We require $q_S(\mathcal{C}) = \mathcal{D}$.

5. **Adding a 1 (respectively, 0)** to the end of each codeword: $S = (1, 2, ..., n, u)$ (respectively $S = (1, 2, ..., n, 0)$). Then $q_S(c) = d$ where $d = (c_1, ..., c_n, 1)$ [respectively $d = (c_1, ..., c_n, 0)$]. This is an extension of our canonical map which added $m - n$ zeros to the end of each codeword; we are now allowed to add a 0 to the end of each codeword, or a 1 to the end of each codeword.

We now define a class of homomorphisms between neural rings which has this property of preserving neuron indicators. This will, it turns out, be exactly the right notion for capturing the maps that we discussed in the introduction.

3.1 Neural ring homomorphisms & proof of Theorem 1.4

The examples listed above are powerful evidence that linear-monomial homomorphisms capture precisely the maps we hoped to be natural code maps. We therefore define *neural ring homomorphisms* and *neural ring isomorphisms* to be the homomorphisms between neural rings which have compatible linear-monomial homomorphisms; this definition formally appears in Section 1.5. In our main result, Theorem 1.4 characterizes the code maps which correspond to neural ring homomorphisms, while Corollary 1.5 does the same for neural ring isomorphisms. We find that code maps which correspond to neural ring homomorphisms are exactly those which are compositions of permutation, dropping the last neuron, adding a codeword, repeating a neuron, or adding a trivial neuron. All are natural code maps for neural data, and the list includes our original motivating maps. We further find that neural ring isomorphisms correspond exactly to permutation code maps, matching our idea that these are the only code maps that completely preserve combinatorial structure. We now prove these results, with help from the following lemmas.

Lemma 3.3. *$\phi : R_{\mathcal{C}} \to R_{\mathcal{D}}$ is a neural ring homomorphism if and only if there is some S such that $q_\phi = q_S\big|_{\mathcal{C}}$.*

Proof. Suppose $\mathcal{C}$ is a code on n neurons and $\mathcal{D}$ a code on m neurons.

If $\phi : R_{\mathcal{D}} \to R_{\mathcal{C}}$ is a neural ring homomorphism then there exists some pure $\tau : R[m] \to R[n]$ compatible with ϕ. Let $\tau = \tau_S$, and let $q_S : \{0,1\}^n \to \{0,1\}^m$ be as above. Then $q_{\tau_S} = q_S$ by Lemma 3.2 and $q_\phi = q_{\tau_S}\big|_{\mathcal{C}}$ by Corollary 2.6 so $q_\phi = q_S\big|_{\mathcal{C}}$. On the other hand, if $q_\phi = q_S\big|_{\mathcal{C}}$ for some S, then ϕ is compatible with τ_S by Corollary 2.6, so as τ_S is pure, ϕ is a neural ring homomorphism. □

Lemma 3.4. *If $q : \mathcal{C} \to \mathcal{D}$ is a restriction of q_T, and $p : \mathcal{D} \to \mathcal{E}$ is a restriction of q_S, then the composition $p \circ q$ is a restriction of q_W for some W.*

Proof of Lemma 3.4. Let $q_T : \mathcal{C} \to \mathcal{D}$ and $q_S : \mathcal{D} \to \mathcal{E}$ be linear-monomial. Suppose $q_T(c) = d$ and $q_S(d) = e$. Let W be defined so $w_i = t_{s_i}$ (where $t_0 = 0$ and $t_u = u$) Then,

$$e_i = \begin{cases} d_j & \text{if } s_i = j \\ 0 & \text{if } s_i = 0 \\ 1 & \text{if } s_i = u \end{cases} = \begin{cases} c_k & \text{or if } s_i = j \text{ and } t_j = k \\ 0 & \text{if } s_i = j \text{ and } t_j = 0 \\ & \text{or if } s_i = 0 \\ 1 & \text{if } s_i = j \text{ and } t_j = u \\ & s_i = u \end{cases}$$

$$= \begin{cases} c_k & \text{if } t_{s_i} = k \\ 0 & \text{if } t_{s_i} = 0 \\ 1 & \text{if } t_{s_i} = u \end{cases} = \begin{cases} c_k & \text{if } w_i = k \\ 0 & \text{if } w_i = 0 \\ 1 & \text{if } w_i = u \end{cases}$$

Thus, $q_S(q_T(c)) = q_W(c)$, and so $q_S \circ q_T$ is defined by W. □

Proof of Theorem 1.4. Lemma 3.3 gives that ϕ is a neural ring homomorphism if and only if $q_\phi = q_S\big|_{\mathcal{C}}$ for some S. So we need only show that $q_\phi = q_S\big|_{\mathcal{C}}$ for some S if and only if q_ϕ is a composition of the listed maps. On one hand, each of the above maps has a related S vector, so by Lemma 3.4 their composition also has one, so if q_ϕ is a composition of the above maps then there is certainly an S such that $q_\phi = q_S\big|_{\mathcal{C}}$.

On the other hand, we show that if $q_\phi = q_S\big|_{\mathcal{C}}$, then we can construct a composition q of the above maps so $q = q_\phi$.

Let $\mathcal{C}$ be a code on n neurons and $\mathcal{D}$ a code on m neurons. Suppose $q_\phi : \mathcal{C} \to \mathcal{D}$ is given by $q_S\big|_{\mathcal{C}}$. code map. For $i = 1, ..., m$, define the function $f_i \in \mathbb{F}_2[x_1, ..., x_n]$ such that $f_i = \begin{cases} x_j & \text{if } s_i = j \\ 1 & \text{if } s_i = u \\ 0 & \text{if } s_i = 0 \end{cases}$

First we define some intermediate codes: let $\mathcal{C}_0 = \mathcal{C}$. For $i = 1, ..., m$, let $\mathcal{C}_i = \{(c_1, ..., c_n, d_1, ..., d_i) \mid c \in \mathcal{C}, d = q(c)\} \subset \{0,1\}^{n+i}$. For $j = 1, ..., n$, let $\mathcal{C}_{m+j} = \{(d_1, ..., d_m, c_1, ..., c_{n-j+1}) \mid c \in \mathcal{C}, d = q(c)\} \subset \{0,1\}^{m+n-j+1}$. Finally, define $\mathcal{C}_{m+n+1} = q(\mathcal{C}) \subset \mathcal{D}$.

Now, for $i = 1, ..., m$, let the code map $q_i : \mathcal{C}_{i-1} \to \mathcal{C}_i$ be defined by $q_i(v) = (v_1, ..., v_{n+i-1}, f_i(v)) \in \mathcal{C}_i$. Note that if $v = (c_1, ..., c_n, d_1, ..., d_{i-1})$, then $f_i(v) = f_i(c)$, as only the first n places matter. Thus, if $v = (c_1, ..., c_n, d_1, ..., d_{i-1})$ with $d = q(c)$, then $q_i(v) = (c_1, ..., c_n, d_1, ..., d_i)$. Neuron by neuron, we add the digits of $q(c)$ on to c. Note that $q_i = q_{S_i}\big|_{\mathcal{C}_{i-1}}$ where $S_i = (1, ..., n+i-1, s_i)$, so q_i is either repeating a neuron, or adding a trivial neuron, depending on whether $s_i = j$, or one of $u, 0$.

Next, take the permutation map given by $\sigma = (n+1, ..., n+m, 1, ..., n)$, so all the newly added neurons are at the beginning and all the originals are at the end. That is, define $q_\sigma : \mathcal{C}_m \to \mathcal{C}_{m+1}$ so if $v = (v_1, ..., v_{n+m})$, then $q_\sigma(v) = (v_{n+1}, ..., v_{n+m}, v_1, ..., v_n)$.

We then delete the neurons $m+1$ through $n+m$ one by one in n code maps. That is, for $j = 1, ..., n$ define $q_{m+j} : \mathcal{C}_{m+j} \to \mathcal{C}_{m+j+1}$ by $q_{m+j}(v) = (v_1, ..., v_{m+n-j})$.

Lastly, if $q(\mathcal{C}) \subsetneq \mathcal{D}$, then add one last inclusion code map $q_a : q(\mathcal{C}) \hookrightarrow \mathcal{D}$ to add the remaining codewords of $\mathcal{D}$.

Thus, given $c = (c_1, ..., c_n)$ with $q(c) = d = (d_1, ..., d_m)$, the first m steps give us $q_m \circ \cdots \circ q_1(c) = (c_1, ..., c_n, d_1, ..., d_m) = x$. The permutation then gives us $q_\sigma(x) = (d_1, ..., d_m, c_1, ..., c_n) = y$, and then we compose $q_{m+n} \circ \cdots \circ q_{m+1}(y) = (d_1, ..., d_n) = d = q(c)$. Finally, if $q(\mathcal{C}) \subsetneq \mathcal{D}$, we do our inclusion map, but as $q_a(d) = d$, the overall composition is a map $\mathcal{C} \to \mathcal{D}$ takes c to $q_S(c) = d$ as desired. At each step, the map we use is from our approved list. □

Proof of Corollary 1.5. Suppose $\tau = \tau_S : R[n] \to R[m]$ is an isomorphism. Then since $R[n]$ and $R[m]$ are finite, they must have the same size. So we must have $2^{2^n} = |R[n]| = |R[m]| = 2^{2^m}$, and thus $n = m$. Thus we can write $\tau : R[n] \to R[n]$.

As τ is an isomorphism, we know $\ker \tau = \{0\}$. So we cannot have $\tau(x_i) = 0$ since then $x_i \in \ker \tau_S$, and similarly we cannot have $\tau(x_i) = 1$, since then $1 - x_i \in \ker \tau$. As τ is linear-monomial, this means that $\tau(x_i) \in \{x_1, ..., x_n\}$ for all $i = 1, ..., n$; as $\tau = \tau_S$ for some S, this means that $s_i \in [n]$ for all i.

If we had $\tau(x_i) = \tau(x_j) = x_k$ for $i \neq j$, then $x_i - x_j \in \ker \tau$, which is a contradiction as $x_i - x_j \neq 0$. So τ induces a bijection on the set of variables $\{x_1, ..., x_n\}$; i.e., S contains each index in $[n]$ exactly once.

Now, consider the corresponding code map q_τ. Let $c \in \{0,1\}^n$, and $q_\tau(c) = d$. We must have $\tau(f)(c) = f(d)$. In particular, we must have $x_j(c) = x_i(d)$, or rather, $c_j = d_i$. So q_τ takes each codeword c to its permutation where $j \to i$ iff $\tau(x_i) = x_j$.

Now, we know that if τ is compatible with ϕ, then ϕ is merely a restriction of the map τ, and so $q_\phi(c) = q_\tau(c)$ for all $c \in \mathcal{C}$. Finally, as ϕ is an isomorphism we know q_ϕ is a bijection, so every codeword in $\mathcal{D}$ is the image of some $c \in \mathcal{C}$; thus, q_ϕ is a permutation map on $\mathcal{C}$, and no codewords are added. □

3.2 Examples of neural ring homomorphisms

Beyond these basic code maps, there are many interesting code maps we can build as compositions. Here are some particular examples.

Deleting an arbitrary set of neurons: While our basic code map above is only defined to delete the final neuron, we can in fact delete any set of neurons. To do this, we first apply a permutation map so the set of neurons to be deleted appear at the end of the world. We then compose as many deletion maps as necessary to remove these neurons and project onto the other coordinates. This composition corresponds to a neural ring homomorphism as it is the composition of basic maps.

Inserting a trivial pattern: By composing permutation maps with maps which add a trivial neuron (1 or 0), we can transform a code to a longer code which has trivial neurons dispersed throughout.

Localization: Localization involves restricting to a piece of the code where a particular activity pattern is present, and looking at the restricted code among the other neurons in this segment. As a code map, this is not well-defined, because there is nowhere to send the codewords which do not exhibit the desired activity pattern. We can, however, consider a kind of reverse localization, where we map the localization back into the code whence it came. This, then, is a composition of adding a trivial pattern (see above), and then an inclusion map. However, not just any inclusion will do - we assume the additional restriction that $\mathcal{D}$ may not contain any codewords exhibiting this activity pattern which aren't in $q(\mathcal{C})$.

4 Neural ring homomorphism and the canonical form

In our previous paper, we focused much of our attention on the *canonical form* for the neural ideal. The canonical form is a distinguished set of minimal generators for the neural ideal which give a condensed yet complete set of combinatorial information about the code. We had previously developed an algorithm for computing the canonical form using the primary decomposition; however, we now present an algorithm which is iterative.

First, we give a few definitions. First note that the neural ideal $I_\mathcal{C} = \{f \in \mathbb{F}_2[x_1, ..., x_n] \mid f(c) = 0 \text{ for all } c \in \mathcal{C}\}$ can be decomposed into two pieces:

$$I_\mathcal{C} = \mathcal{B} + J_\mathcal{C}$$

where $\mathcal{B} = \langle x_i^2 - x_i \mid i \in [n] \rangle$ and $J_\mathcal{C} = \langle \prod_{v_i=1} x_i \prod_{v_i=0}(1 - x_i) \mid v \in \{0,1\}^n \backslash \mathcal{C} \rangle$. Polynomials of the form $\prod_{i\in\sigma} x_i \prod_{j\in\tau}(1-x_j)$, where $\sigma \cap \tau = \emptyset$, are referred to as *pseudo-monomials*, and the ideal $J_\mathcal{C}$, termed the *neural ideal*, can be generated by a set of minimal pseudo-monomials which we call the canonical form. For more on this canonical form $CF(J_\mathcal{C})$ see [4].

Example. Consider the code $\mathcal{C} = \{000, 100, 010, 001, 110, 011\}$. $J_\mathcal{C}$ is generated by $\langle x_1x_2x_3, x_1(1-x_2)x_3\rangle$, but the canonical form is given by $CF(J_\mathcal{C}) = \{x_1x_3\}$, which also generates $J_\mathcal{C}$.

For the most part, it is simple to describe the effects of our basic code maps on the canonical form. Two proofs (repeating a neuron, adding a codeword) are presented in full detail; the others follow a similar structure and the intuition is briefly sketched. Throughout, $\mathcal{C}$ is a code on n neurons with canonical form $CF(J_\mathcal{C})$, and $z_i \in \{x_i, 1-x_i\}$ represents one of the two possible non constant linear terms involving x_i. Because $J_\mathcal{C}$ is an ideal of $\mathbb{F}_2[x_1, ..., x_n]$ we revert back to polynomial notation entirely in this section.

1. **Permutation:** as this map simply permutes the labels on the variables, the canonical form stays nearly the same, but the labels are permuted using the reverse permutation σ^{-1}. That is, let $\mathcal{D}$ be the code obtained by applying the permutation $\sigma \in \mathcal{S}_n$ to $\mathcal{C}$. Then $f = z_{i_1} \cdots z_{i_k} \in CF(J_\mathcal{C})$ if and only if $f_\sigma = z_{\sigma^{-1}(i_1)} \cdots z_{\sigma^{-1}(i_k)} \in CF(J_\mathcal{D})$. That is, all relationships amongst the variables stay the same; only the labels change to reflect the permutation.

2. **Adding a codeword:** This is the most complex operation on the canonical form. The algorithmic process for obtaining $CF(\mathcal{C} \cup \{v\})$ from $CF(\mathcal{C})$ is described in detail in the following section, and forms the basis for our iterative algorithm. Essentially, we retain those generators which still apply under the new codeword; those which do not are adjusted by multiplication with an appropriate linear term.

3. **Deleting a neuron:** Let $\mathcal{C}$ be a code on n neurons, and $\mathcal{D}$ the code on $n-1$ neurons obtained by deleting the nth neuron. Then $CF(J_\mathcal{D}) = CF(J_\mathcal{C}) \setminus \{f \mid f = gz_n, g \text{ a pseudo-monomial}\}$. That is, we simply remove all pseudo-monomials which involved the variable x_n; however, the relationships amongst the others are unchanged. We never add new relations, because any relation that holds without neuron n held before, and would have appeared.

4. **Adding a trivial neuron which is always 1 (always 0):** Let $\mathcal{D}$ be the code on $n+1$ neurons obtained by adding a 1 (respectively 0) to the end of each codeword in $\mathcal{C}$. Then $CF(J_\mathcal{D}) = CF(J_\mathcal{C}) \cup \{1 - x_{n+1}\}$ (respectively $CF(J_\mathcal{C}) \cup \{x_{n+1}\}$). Here, we add a single generator to indicate that this new neuron is uniformly 1 (or 0) and the relationships amongst the others are unchanged.

5. **Adding a new neuron which repeats another neuron:** Let $\mathcal{D}$ be the code on $n+1$ neurons obtained from $\mathcal{C}$ by adding a new neuron which repeats neuron i for all codewords.

   Let $F = \{f \in CF(J_\mathcal{C}) \mid f = z_i \cdot g \text{ for } g \text{ a pseudo monomial}\}$; let H be the set formed by replacing x_i with x_{n+1} for all $f \in F$. Then in most cases, $CF(J_\mathcal{D}) = CF(J_\mathcal{C}) \cup \{x_i(1 - x_{n+1}), x_{n+1}(1 - x_i)\} \cup H$. The exception is when $z_i \in CF(J_\mathcal{C})$; then $CF(J_\mathcal{D})$ is simply $CF(J_\mathcal{C}) \cup H = CF(J_\mathcal{C}) \cup \{z_{n+1}\}$.

   Proof. In the first case, suppose $z_n \in CF(J_\mathcal{C})$. Then we know the nth bit is either always 1 or always 0, so by repeating it we are really adding a new neuron which is always 1 or always 0. We can then follow the changes described above, which means we merely add the appropriate z_{n+1} term.

   In the second case, suppose $z_n \notin CF(J_\mathcal{C})$. First, note that $CF(J_\mathcal{C}) \subseteq J_\mathcal{D}$, as a polynomial in $CF(J_\mathcal{C})$ involves the first n variables only, and the first n bits of any codeword in $\mathcal{D}$ are a copy of a word in $\mathcal{C}$. Note also that any word in $\mathcal{D}$ has the same $n+1$st digit as nth digit, so if we take a polynomial from $CF(J_\mathcal{C})$ which has a z_n term and replace it with the corresponding

z_{n+1} term it will still be 0 on all elements of $\mathcal{D}$; thus $H \subseteq J_{\mathcal{D}}$. Lastly, since each polynomial has the same n and $n+1$ digits, we have $x_n(1-x_{n+1})$ and $x_{n+1}(1-x_n) \in J_{\mathcal{D}}$. Note also that any polynomial in $J_{\mathcal{D}}$ which doesn't involve x_{n+1} is also a polynomial in $J_{\mathcal{C}}$.

Suppose f is a minimal pseudo-monomial in $CF(J_{\mathcal{D}})$. Then we have three cases:

(a) z_n and z_{n+1} both divide f. If both linear terms are the same, then removing one of them also must give something in $J_{\mathcal{D}}$ (because those two digits are always the same) so they're not really minimal. If they're different, then f is a multiple of either $x_n(1-x_{n+1})$ or $x_{n+1}(1-x_n)$ and since these are minimal it must be one of these.

(b) $z_{n+1} \big| f$, but z_n does not. Then, changing x_{n+1} to x_n gives a pseudo-monomial g which appears in $J_{\mathcal{C}}$. It must be minimal, or else its divisor would be a divisor of f (after changing x_{n+1} back to x_n if necessary). Thus $g \in CF(J_{\mathcal{C}})$ and hence $f \in H$.

(c) $z_{n+1} \not| f$. Then $f \in J_{\mathcal{C}}$. But then it must also be minimal in $J_{\mathcal{C}}$, or else its divisor would also be in $J_{\mathcal{D}}$ which would contradict the stated minimality. So $f \in CF(J_{\mathcal{C}})$.

Together, this shows that $CF(J_{\mathcal{D}}) \subseteq CF(J_{\mathcal{C}}) \cup H \cup \{x_n(1-x_{n+1}), x_{n+1}(1-x_n)\}$.

By similar arguments, we get reverse containment. If $f \in CF(J_{\mathcal{C}})$, then it is minimal in $J_{\mathcal{D}}$ as well, as any divisor of it would also be in $J_{\mathcal{C}}$. If $f \in H$, then it must be minimal, because the polynomial g obtained by switching x_n back in for x_{n+1} was minimal, so $f \in CF(J_{\mathcal{D}})$. And both $x_n(1-x_{n+1})$ and $x_{n+1}(1-x_n)$ are minimal, as we don't have $z_n \in J_{\mathcal{C}}$ so the nth bit sometimes takes value 0 and sometimes value 1. Thus we can't have z_{n+1} in $J_{\mathcal{D}}$ because the $n+1$st bit also varies. □

4.1 Iterative algorithm for computing $CF(J_{\mathcal{C}})$

We now describe the algorithm for computing $CF(J_{\mathcal{C}})$, starting with the canonical form for a single codeword and iteratively applying the map which adds a single codeword until the whole code $\mathcal{C}$ is present. The code for this algorithm is available online [7].

Input: $\mathcal{C} \subset \{0,1\}^n$, a code consisting of m codewords. Order codewords arbitrarily $c^1, ..., c^m$.
Output: The canonical form for $J_{\mathcal{C}}$.

Description: Let $\mathcal{C}^k$ be the restriction of $\mathcal{C}$ to the first k codewords. $CF(J_{\mathcal{C}^1})$ is simple to describe:

$$CF(J_{\mathcal{C}^1}) = \{x_i - c_i^1 \mid i \in [n]\}.$$

We create $CF(J_{\mathcal{C}^k})$ from $CF(J_{\mathcal{C}^{k-1}})$ for $2 \leq k \leq m$, by the process described in Algorithm 1, described below. After $m-1$ iterations, we obtain $CF(J_{\mathcal{C}^m}) = CF(J_{\mathcal{C}})$. If desired, one also could store the successive iterations and output the canonical form for all the nested codes.
Algorithm 1 describes the process to compute $CF(J_{\mathcal{C}^k})$ from $CF(J_{\mathcal{C}^{k-1}})$ and c^k. To do this, we note that $\mathcal{C}^k = \mathcal{C}^{k-1} \cup \{c^k\}$; hence we can use Lemma 4.2, which gives a process to find the canonical form for the union of two codes if we know the canonical forms of each respective ideal. We have $CF(J_{\mathcal{C}^{k-1}}) = \{f_1, ..., f_r\}$, and $CF(J_{\{c^k\}}) = \{x_i - c_i^k \mid i \in [n]\}$. By the lemma, we obtain $CF(J_{\mathcal{C}^k})$ by taking all possible products $f_i(x_j - c_j^k)$ and reducing the 0s and the multiples.

However, since one of our canonical forms consists of only linear terms, we can take a few shortcuts along the way, as described in the following algorithm.

Algorithm 1: Computing $CF(J_{\mathcal{C}^k})$ from $CF(J_{\mathcal{C}^{k-1}})$ and c^k:

Input: $CF(J_{\mathcal{C}^{k-1}}) = \{f_1, ..., f_r\}$ and $CF(J_{\{c^k\}}) = \{x_i - c_i^k \mid i \in [n]\}$.
Output: $CF(J_{\mathcal{C}^k})$, where $\mathcal{C}^k = \mathcal{C}^{k-1} \cup \{c^k\}$.

Step 1: Let $M = \{f_i \in CF(J_{\mathcal{C}^{k-1}}) \mid f_i(c^k) = 0\}$, and $N = CF(J_{\mathcal{C}^{k-1}}) \backslash M$. Initialize a third set, L, as an empty set.

Step 2: There are $n|N|$ possible products of polynomials f_j in N and linear terms $x_i - c_i^k$ in $CF(J_{\{c^k\}})$. Order them arbitrarily.

Step 3: Consider these products one by one. For each product $g = f_j(x_i - c_i^k)$, there are 3 possibilities:

a) If $x_i - c_i^k - 1$ divides f_j, discard g and move on to the next product.

b) If a) does not hold, but g is a multiple of an element of M, discard g and move on to the next product.

c) If neither a) nor b) holds, add g to L and move on to the next product.

Repeat until all products have been considered.

Step 4: Output $M \cup L$.

Proposition 4.1. *The process described in Algorithm 1 gives the canonical form for $CF(J_{\mathcal{C}}^k)$ from $CF(J_{\mathcal{C}}^{k-1})$ and $CF(J_{c^k})$.*

4.2 Proof of Algorithm 1

Throughout, $\mathcal{C}$ and $\mathcal{D}$ are neural codes on the same number of neurons; so, $\mathcal{C}, \mathcal{D} \subseteq \{0,1\}^n$. All polynomials are elements of $\mathbb{F}_2[x_1, ..., x_n]$. We will use the following useful facts about $J_{\mathcal{C}}$:

- If f is a pseudo-monomial and $f(c) = 0$ for all $c \in \mathcal{C}$, then $f \in J_{\mathcal{C}}$.
- If $f \in J_{\mathcal{C}}$ is a pseudo-monomial, then $f = gh$ for some $h \in CF(J_{\mathcal{C}})$ and g a pseudo-monomial.
- If $\mathcal{C} \subseteq \mathcal{D}$, then $J_{\mathcal{C}} \supseteq J_{\mathcal{D}}$.

We also make use of the following definitions: a monomial x^α is *square-free* if $\alpha_i \in \{1, 0\}$ for all $i = 1, ..., n$. A polynomial is *square-free* if it can be written as the sum of square-free monomials. For example: $x_1x_2 + x_4 + x_1x_3x_2$ is square-free. There is a unique square free representative of every equivalence class of $\mathbb{F}_2[x_1, ..., x_n]/\mathcal{B}$. To see this, recall that $\mathbb{F}_2[x_1, ..., x_n]/\mathcal{B}$ is the set of functions from $\{0,1\}^n \to \{0,1\}$. There are 2^{2^n} such functions; thus there are 2^{2^n} equivalence classes. Any polynomial representative of one of these classes can be reduced to a square-free polynomial representative by applying the relations $x_i = x_i^2$ from $\mathcal{B}$. As there are 2^n square-free monomials in $\mathbb{F}_2[x_1, ..., x_n]$, there are exactly 2^{2^n} square-free polynomials total; thus each equivalence class must have a unique representative.

For $h \in \mathbb{F}_2[x_1, .., x_n]$, let h_R denote the unique square-free representative of the equivalence class of h. This is similar to choosing the equivalence class of h mod $\mathcal{B}$, but we select a very particular representative of that class.

Then, for $CF(J_{\mathcal{C}}) = \{f_1, ..., f_r\}$ and $CF(J_{\mathcal{D}}) = \{g_1, ..., g_s\}$, we define the set of *reduced products*

$$P(\mathcal{C}, \mathcal{D}) \stackrel{\text{def}}{=} \{(f_i g_j)_R \,|\, i \in [r], j \in [s]\}$$

and the *minimal reduced products* as

$$MP(\mathcal{C}, \mathcal{D}) \stackrel{\text{def}}{=} \{h \in P(\mathcal{C}, \mathcal{D}) \,|\, h \neq 0 \text{ and } h \neq fg \text{ for any } f \in P(\mathcal{C}, \mathcal{D}), g \neq 1\}.$$

Lemma 4.2. *If $\mathcal{C}, \mathcal{D} \subset \{0,1\}^n$, then the canonical form of their union is given by the set of reduced products from their canonical forms: $CF(J_{\mathcal{C}\cup\mathcal{D}}) = MP(\mathcal{C}, \mathcal{D})_R$.*

Proof. Let $h = (f_i g_j)_R \in MP(\mathcal{C}, \mathcal{D})$. Then h is a pseudo-monomial, and note that $h \in J_{\mathcal{C}}$ as it is a multiple of f_i, and $h \in J_{\mathcal{D}}$ as it is a multiple of g_j. Thus $h(c) = 0$ for all $c \in \mathcal{C} \cup \mathcal{D}$, so $h \in J_{\mathcal{C}\cup\mathcal{D}}$.

Now, let $h \in CF(J_{\mathcal{C}\cup\mathcal{D}})$. Then as $J_{\mathcal{C}\cup\mathcal{D}} \subset J_{\mathcal{C}}$, there is some $f_i \in CF(J_{\mathcal{C}})$ so that $h = h_1 f_i$, and likewise there is some $g_j \in CF(J_{\mathcal{D}})$ so $h = h_2 g_j$ where h_1, h_2 are also pseudo monomials. Thus h is a multiple of $(f_i g_j)_R$ and hence is a multiple of some element of $MP(\mathcal{C}, \mathcal{D})$. But as every element of $MP(\mathcal{C}, \mathcal{D})$ is an element of $J_{\mathcal{C}\cup\mathcal{D}}$, and $h \in CF(J_{\mathcal{C}\cup\mathcal{D}})$, this means h must actually be in $MP(\mathcal{C}, \mathcal{D})$.

Thus, $CF(J_{\mathcal{C}\cup\mathcal{D}}) \subseteq MP(\mathcal{C}, \mathcal{D})$.

Finally, to see that $MP(\mathcal{C}, \mathcal{D}) \subseteq CF(J_{\mathcal{C}\cup\mathcal{D}})$, if $h \in MP(\mathcal{C}, \mathcal{D})$, then h is in $J_{\mathcal{C}\cup\mathcal{D}}$. It is thus the multiple of some $f \in CF(J_{\mathcal{C}\cup\mathcal{D}})$. But we have shown that $f \in MP(\mathcal{C}, \mathcal{D})$, which contains no multiples. So $h = f$ is in $CF(J_{\mathcal{C}\cup\mathcal{D}})$. □

Proof. (Proof of Proposition 4.1) Since we are considering products between the canonical form for a code $\mathcal{C}^{k-1}$ and that for a single codeword c^k, we will dispense with the excess indices and use $\mathcal{C}$ and c, respectively. Note that if $c \in \mathcal{C}$, then $M = CF(J_{\mathcal{C}})$, so the algorithm ends immediately and outputs $CF(J_{\mathcal{C}})$; we will generally assume $c \notin \mathcal{C}$. To show that the algorithm produces the correct canonical form, we apply Lemma 4.2, so it suffices to show that the set $M \cup L$ is exactly $MP(\mathcal{C}, \{c\})$. This requires that all products are considered, and that we remove exactly those which are multiples or zeros.

To see that all products are considered: Let $g \in M$. Since $g(c) = 0$, we know $(g(x_i - c_i))_R = g$ for at least one i. So $g \in MP(\mathcal{C}, \{c\})$. Any other product $(gz_j)_R$ will either be g, or will be a multiple of g, and hence will not appear in $MP(\mathcal{C}, \{c\})$. Thus, all products of linear terms with elements of M are considered, and all multiples or zeros are removed. It is impossible for elements of M to be multiples of one another, as $M \subset CF(\mathcal{C})$.

We also consider all products of elements of N with the linear elements of $CF(J_c)$. We discard them if their reduction would be 0 (case (a)) or if they are a multiple of anything in M (case (b)). If neither holds, we keep them in L. So it only remains to show that no element of L can be a multiple of any other element in L, and no element of M can be a multiple of anything in L, and thus that we have taken care of every possibility.

First, no element of M may be a multiple of an element of L, since if $g \in M$, $fz_i \in L$, and $fz_i \cdot p = g$ for some pseudo-monomial p, then $f \big| g$. But this is impossible as f, g are both in $CF(J_{\mathcal{C}})$.

Now, suppose $fz_i = h \cdot gz_j$ for $f, g \in CF(J_{\mathcal{C}})$ and $fz_i, gz_j \in L$, and h a nontrivial pseudo-monomial. Then as $f \not| g$ and $g \not| f$, we have $i \neq j$, and so $z_j \big| f$. But this means $fz_j = f$ and therefore $f \in M$, which is a contradiction. So no elements of L may be multiples of one another. □

References

[1] L. Osborne, S. Palmer, S. Lisberger, and W. Bialek. The neural basis for combinatorial coding in a cortical population response. *Journal of Neuroscience*, 28(50):13522–13531, 2008.

[2] C. Curto, V. Itskov, K. Morrison, Z. Roth, and J. Walker. Combinatorial neural codes from a mathematical coding theory perspective. *Neural computation*, 25(7):1891–1925, 2013.

[3] C. Curto and V. Itskov. Cell groups reveal structure of stimulus space. *PLoS Computational Biology*, 4(10), 2008.

[4] C. Curto, V. Itskov, A. Veliz-Cuba, and N. Youngs. The neural ring : an algebraic tool for analyzing the intrinsic structure of neural codes. *Bulletin of Mathematical Biology*, 75(9), 2013.

[5] C. Curto, E. Gross, J. Jeffries, K. Morrison, M. Omar, Z. Rosen, A. Shiu, and N. Youngs. What makes a neural code convex? *arXiv:1508.00150; Submitted*, 2015.

[6] D. Cox, J. Little, and D. O'Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, second edition, 1997. An introduction to computational algebraic geometry and commutative algebra.

[7] N. Youngs. Neural ideal: a matlab package for computing canonical forms. `https://github.com/nebneuron/neural-ideal`, 2015.